Resource-Constrained Adaptive Search for Sparse Multi-Class Targets with Varying Importance

Gregory E. Newstadt, *Member, IEEE,* Beipeng Mu, *Member, IEEE,* Dennis Wei, *Member, IEEE,* Jonathan P. How *Senior Member, IEEE*, and Alfred O. Hero III, *Fellow, IEEE*

arXiv:1409.7808v1 [cs.IT] 27 Sep 2014

Abstract

In sparse target inference problems it has been shown that significant gains can be achieved by adaptive sensing using convex criteria. We generalize previous work on adaptive sensing to (a) include multiple classes of targets with different levels of importance and (b) accommodate multiple sensor models. New optimization policies are developed to allocate a limited resource budget to simultaneously locate, classify and estimate a sparse number of targets embedded in a large space. Upper and lower bounds on the performance of the proposed policies are derived by analyzing a baseline policy, which allocates resources uniformly across the scene, and an oracle policy which has a priori knowledge of the target locations/classes. These bounds quantify analytically the potential benefit of adaptive sensing as a function of target frequency and importance. Numerical results indicate that the proposed policies perform close to the oracle bound when signal quality is sufficiently high (e.g. performance within 3 dB for SNR above 15 dB). Moreover, the proposed policies improve on previous policies in terms of reducing estimation error, reducing misclassification probability, and increasing expected return. To account for sensors with different levels of agility, three sensor models are considered: global adaptive (GA), which can allocate different amounts of resource to each location in the space; global uniform (GU), which can

Gregory Newstadt is with Google Pittsburgh, Pittsburgh, PA 15206, E-mail: gregoryn@google.com

Beipeng Mu and Jonathan How are with Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, E-mail: ({mubp},{jhow})@mit.edu

Dennis Wei is with the Thomas J. Watson Research Center, IBM Research, Yorktown Heights, NY 10598, E-mail: dwei@us.ibm.com.

Alfred Hero is with the Dept. of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, E-mail: hero@umich.edu.

The research in this paper was partially supported by Army Research Office MURI grant number W911NF-11-1-0391.

allocate resources uniformly across the scene; and local adaptive (LA), which can allocate fixed units to a subset of locations. Policies that use a mixture of GU and LA sensors are shown to perform similarly to those that use GA sensors while being more easily implementable.

I. Introduction

This work considers localization, classification, and estimation of targets from observations taken sequentially and adaptively. In particular, we focus on the regime where (a) targets are sparse compared to the size of the scene, and (b) some of the targets have higher importance than others. For example, in search-and-rescue missions, detection of survivors has significantly higher mission importance than detection of other features in the environment. Similarly, a radar operator may be more interested in detecting/tracking a tank rather than a car, though both might sparsely populate the scene.

Viewing the targets as sparse signals, adaptive localization and estimation use past observations to shape future measurements of the scene [1]–[10], which can result in stronger signal-to-noise ratios. The performance gain occurs by adaptively focusing the majority of sensing resources only on the positions that contain targets. Applications where adaptive sensing has been used include image acquisition/compression, spectrum sensing, agile radars, and medical imaging [4,6,8,11]–[13].

Previous work on adaptive sensing for sparse targets [3]–[14] considered only the two-class detection problem, where the target is either absent or present. In many applications, such as surveillance and search-and-rescue missions, targets may have different classes with varying importance to the mission. In this setting, detection-based methods may waste resources on unimportant targets and thus suffer performance losses for the important ones. This work explicitly accounts for multi-class targets with different mission importance.

Previous work in adaptive sensing [3,6,8] also assumed the availability of an agile sensor that can allocate sensing resources to any combination of locations in the scene at the same time and with potentially different effort (i.e. dwell time or energy). In some cases we may only have access to sensors with restricted agility in terms of allocating different quantities of resources to different locations. For example, a global sensor (e.g. in wide-area surveillance) may only be able to allocate resources uniformly across the scene; a local sensor (e.g. on an unmanned aerial vehicle) may only be able to allocate resources to a small subset of the scene. This paper provides resource allocation policies for when a globally adaptive sensor is available, as well as when the agility of the sensor is limited. Moreover, the differences in performance of these policies are analyzed and the regimes where each of the policies perform well are identified.

There is relatively limited work for planning with multiple types of sensors and tasking agents, particularly when the size of the scene is large. Previous work [15,16] proposed a planning algorithm for a team of heterogeneous sensors with the goal of maximizing mission importance. This algorithm showed that tasking agents achieved significantly better performance gains when target exploration included mission importance in the planning stage, as compared to an exploration policy that only depended on target uncertainty. This resource allocation problem included many additional parameters (e.g. timing and delays), but was limited to the situation where the number and location of targets were assumed to be known a priori. In this work, we relax this assumption and simultaneously enumerate and localize targets along with sensor planning.

We provide a Bayesian formulation and objective function that generalize our previous work [3,8] to include targets with different mission importance. Subsequently, upper bounds on the performance of adaptive resource allocation policies are derived through analysis of a full oracle policy with complete knowledge of the target locations and mission importances and a location-only oracle with knowledge only of target locations. Similarly, a lower bound on performance is derived from analysis of a baseline policy which uniformly allocates resources everywhere in the scene. These bounds indicate that there are regimes in which incorporating mission importance can lead to significant gains over both the baseline policy, as well as adaptive policies which do not incorporate mission importance. Moreover, the bounds yield a simple method for mathematically predicting the benefit of adaptive sensing as a function of the system and target parameters.

This paper also provides resource allocation policies that are implementable (in contrast to the oracle policies) and approximately optimize the objective function in multiple scenarios, depending on the agility of the available sensors. The performance of these policies is compared numerically to baseline policies, oracle policies and previously proposed policies, and the regimes (i.e., conditions on model parameters) in which the new policies perform well are identified. Results indicate that the proposed policies perform significantly better than policies which do not include mission importance. Moreover, in the high-resource regime, the proposed approach performs nearly as well as the oracle policies. In scenarios where a globally adaptive sensor is not available, we further identify the tradeoffs between performance and sensor complexity (e.g., number of local sensors and/or inclusion of a global uniform sensor). Finally, while the objective function is chosen for its convexity and ease of optimization, it is shown numerically that our proposed approach performs well over many performance metrics, including reducing mean squared estimation error and misclassification probability.

The rest of the paper is organized as follows: Section II presents the target and observation models,

as well as the objective function; Section III derives and compares the performance bounds; Section IV provides approximate optimization methods for multiple sensing scenarios; Section V compares performance of the proposed policies; and the paper concludes in Section VI.

II. MODEL

Consider a scene with N discrete locations, which contains a small number of targets of interest, for example, vehicles in wide area imagery. The i-th location has an unknown class label $C_i \in \mathcal{C} = \{1, 2, \dots\}$, where $C_i = 1$ indicates the absence of a target. Define $\{p_c\}_{c\in\mathcal{C}}$ as the prior distribution of the target class at a location, where $p_c = \Pr(C_i = c)$ and $\sum_{c\in\mathcal{C}} p_c = 1$. Thus, the few-target (i.e. sparsity) assumption can be restated as $p_1 \approx 1$. The prior distribution is assumed to be known. Moreover, this formulation can be easily extended so that the prior probabilities also depend on location.

Correct classification of a location to class c yields reward $h(c)$ to the mission (hereafter referred to as the mission importance). Here, $h : \mathcal{C} \to [0, \infty)$ is a known reward function that maps the target class to its mission importance. Moreover, it is assumed that the value of the no-target class is always zero, so that $h(1) = 0$. An example of a 3-class problem includes the no-target class ($c = 1$), a low-value target class ($c = 2$, e.g. a car) and high-value target class ($c = 3$, e.g. a tank) where

$$h(c) = \begin{cases} 0, & c = 1 \\ 1, & c = 2 \\ 10, & c = 3 \end{cases}. \tag{1}$$

Without loss of generality, $h(c)$ is also assumed to be monotonically increasing in c.

Associated with each location i is an amplitude X_i, called the signal and modeled as follows:

$$X_i | C_i = c \sim \begin{cases} \text{Normal}(\mu_c, \sigma_c^2), & c > 1 \\ 0, & c = 1 \end{cases} \tag{2}$$

In words, conditioned on the event $C_i = c$ for $c > 1$, X_i is distributed as a Gaussian random variable with known mean μ_c and variance σ_c^2, while X_i is equal to zero with probability one if $c = 1$.

Observations of the signals are collected over T stages with resource allocations $\lambda_i(t)$, which are a function of both location i and stage index $t = 0, \cdots, T-1$. The resource $\lambda_i(t)$ can depend on observations up to and including stage t. Examples of resource $\lambda_i(t)$ include integration time, number of samples, or transmitted power. Given $\lambda_i(t-1) > 0$, the next observation $y_i(t)$ takes the form:

$$y_i(t) = X_i + \frac{n_i(t)}{\sqrt{\lambda_i(t-1)}},\ i = 1, \dots, N,\ t = 1, \dots, T \tag{3}$$

where $\{n_i(t)\}_{i,t}$ is i.i.d. zero-mean Gaussian noise with variance ν^2. If $\lambda_i(t-1)=0$, no observation is taken. A key point of this model is that the observation quality increases with sensing effort (i.e., $\lambda_i(t-1)$). The resource allocations are constrained with the following total resource budget:

$$\sum_{t=0}^{T-1}\sum_{i=1}^{N}\lambda_i(t)=\sum_{t=0}^{T-1}\Lambda(t)=\Lambda, \tag{4}$$

where $\Lambda(t)$ are the per-stage budgets. For convenience, we write $\boldsymbol{\lambda}(t)=[\lambda_1(t)\ \lambda_2(t)\dots\lambda_N(t)]^T$, $\boldsymbol{y}(t)=[y_1(t)\ y_2(t)\dots y_N(t)]^T$ (similarly for other indexed quantities) and define $\boldsymbol{Y}(t)=\{\boldsymbol{y}(1),\boldsymbol{y}(2),\dots,\boldsymbol{y}(t)\}$. The sequence of effort allocations $\boldsymbol{\lambda}=\{\boldsymbol{\lambda}(t)\}_{t=0}^{T-1}$ is called the allocation policy, where $\boldsymbol{\lambda}(t)$ is a mapping from $\boldsymbol{Y}(t)$ to $[0,\Lambda(t)]^N$.

A straightforward multi-class extension of results in [8] yields the following posterior model for targets and their associated signals. Conditioned on the measurements up until time t and the class C_i, we have

$$X_i\big|\boldsymbol{Y}(t),C_i=c\sim\text{Normal}(\hat{X}_i^{(c)}(t),(\sigma_i^{(c)}(t))^2), \tag{5}$$

where the posterior mean (conditional mean estimator) and variance are defined as

$$\hat{X}_i^{(c)}(t)=\mathbb{E}\left[X_i\big|\boldsymbol{Y}(t),C_i=c\right] \tag{6}$$

$$(\sigma_i^{(c)}(t))^2=\text{var}\left[X_i\big|\boldsymbol{Y}(t),C_i=c\right] \tag{7}$$

with $\hat{X}_i^{(c)}(0)=\mu_c$ and $(\sigma_i^{(c)}(0))^2=\sigma_c^2$. The posterior probability for a target at location i having class c is

$$p_i^{(c)}(t)=\Pr(C_i=c|\boldsymbol{Y}(t)) \tag{8}$$

with $p_i^{(c)}(0)=p_c$. When the observation $y_i(t)$ is taken, (6)-(8) satisfy the update equations

$$(\sigma_i^{(c)}(t))^2=\nu^2\left[\frac{\nu^2}{(\sigma_i^{(c)}(t-1))^2}+\lambda_i(t-1)\right]^{-1} \tag{9}$$

$$\hat{X}_i^{(c)}(t)=(\sigma_i^{(c)}(t))^2\left(\frac{\hat{X}_i^{(c)}(t-1)}{(\sigma_i^{(c)}(t-1))^2}+\frac{y_i(t)\lambda_i(t-1)}{\nu^2}\right) \tag{10}$$

$$p_i^{(c)}(t)=\frac{p_i^{(c)}(t-1)f(y_i(t)|\boldsymbol{Y}(t-1),C_i=c)}{\sum_{c'\in\mathcal{C}}p_i^{(c')}(t-1)f(y_i(t)|\boldsymbol{Y}(t-1),C_i=c')} \tag{11}$$

where $y_i(t+1)|\boldsymbol{Y}(t),C_i=c$

$$\sim\begin{cases}\text{Normal}\left(0,\nu^2/\lambda_i(t)\right) & c=1,\\ \text{Normal}\left(\hat{X}_i^{(c)}(t),(\sigma_i^{(c)}(t))^2+\nu^2/\lambda_i(t)\right), & c>1,\end{cases} \tag{12}$$

If $y_i(t)$ is not taken, the posterior parameters do not change.

A. Objective function for resource allocations

We consider a generalization of the resource allocation objective functions in [3,8]. In analogy with [3,8], define $I_i^{(c)}$ to be the following indicator function

$$I_i^{(c)} = \begin{cases} 1, & C_i = c \\ 0, & C_i \neq c \end{cases}. \tag{13}$$

The proposed objective function can then be expressed as

$$J_T(\boldsymbol{\lambda}) = \mathbb{E}_{\boldsymbol{Y}(T),\boldsymbol{C},\boldsymbol{X}} \left[\sum_{i=1}^{N} \sum_{c=2}^{|\mathcal{C}|} I_i^{(c)} h(c) \left(X_i - \hat{X}_i^{(c)}(T) \right)^2 \right], \tag{14}$$

where $\mathbb{E}_{\boldsymbol{Y}(T),\boldsymbol{C},\boldsymbol{X}}[\cdot]$ is the expectation operator over $\boldsymbol{Y}(T), \boldsymbol{C}$, and $\boldsymbol{X}$. The inner sum corresponds to the mean-squared error (MSE) in estimating X_i given all observations $\boldsymbol{Y}(T)$ and class C_i, weighted by mission importance $h(C_i)$. Note that we have already substituted the posterior mean estimator $\hat{X}_i^{(c)}(T)$ in (14), which minimizes the MSE.

The two-class problem considered in [3,8] is a special case of (14) with $\mathcal{C} = \{1, 2\}$, $h(1) = 0$ and $h(2) = 1$. These previous works derived optimal resource allocation policies that minimize (14) when there is only a single target class with nonzero mission importance. The generalization to multiple nonzero target classes in (14) penalizes the policy for incorrect target classification in addition to incorrect detection.

To explicitly reflect the fact that the target classes are random and not known exactly, we may rewrite (14) in terms of the posterior class probabilities $p_i^{(c)}(t)$ instead of the indicators $I_i^{(c)}$, as shown below.

Proposition 1 The objective function (14) is equivalent to

$$\begin{aligned} J_T(\boldsymbol{\lambda}) &= \nu^2 \mathbb{E}_{\boldsymbol{Y}(T-1)} \left[\sum_{i=1}^{N} \sum_{c=2}^{|\mathcal{C}|} \frac{p_i^{(c)}(T-1) h(c)}{\nu^2/(\sigma_i^{(c)}(T-1))^2 + \lambda_i(T-1)} \right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{Y}(T-1)} \left[\sum_{i=1}^{N} \sum_{c=2}^{|\mathcal{C}|} \frac{p_i^{(c)}(T-1) h(c)}{\nu^2/\sigma_c^2 + \overline{\lambda_i}} \right], \end{aligned} \tag{15}$$

where $\overline{\lambda_i} = \sum_{t=0}^{T-1} \lambda_i(t)$.

Proof: The proof is given in Appendix IV. ■

The alternative form of the objective function in (15) also shows its explicit dependence on the allocated resources $\boldsymbol{\lambda} = \{\boldsymbol{\lambda}(t)\}_{t=0}^{T-1}$, which will be useful in the sequel.

While it is apparent from (14) that the objective function explicitly penalizes estimation errors, classification errors are implicitly penalized as well. More specifically, (14) encourages the allocation of more resources to reduce the MSE in locations where targets have high value $h(C_i)$. However, in order to selectively concentrate resources in this way, targets must also be correctly classified. Proposition 2 confirms the intuition that allocating resource to a location reduces the corresponding mis-classification error. We consider a modified Maximum Posterior (MAP) class estimator:

$$\hat{C}_i(t) = \begin{cases} \arg\max_{c\in\mathcal{C}} \ p_i^{(c)}(t) & \text{when } y_i(t) \geq \epsilon \\ 1 & \text{when } y_i(t) < \epsilon \end{cases} \tag{16}$$

where ϵ is a user-chosen threshold. Define misclassification probability of location i with true class C_i, given measurements up until stage t as

$$q_i(t) = \Pr\left(\hat{C}_i(t) \neq C_i \big| \boldsymbol{Y}(t-1)\right) \tag{17}$$

We define $\sigma_0^2 \stackrel{\triangle}{=} \max_{c>1} \sigma_c^2$ and introduce the following assumption of equal prior target variances:

Assumption 1 $\sigma_c^2 = \sigma_0^2$ for all $c > 1$.

This assumption leads to upper bounds on both the mis-classification error as well as the objective function (15) since target variances smaller than the maximum would make estimation and classification easier.

Proposition 2 Under Assumption 1, when $\hat{X}_i^{(c)}(t-1) > 0$ for $c > 1$ and $\epsilon = 0$, $q_i(t)$ is monotonically decreasing with $\lambda_i(t-1)$.

Assumption 1 and the condition $\hat{X}_i^{(c)}(t-1) > 0$ for $c > 1$ simplify the proof of Proposition 2 given in Appendix I. If $\mu_c > 0$ for $c > 1$ and the sensor noise variance ν^2 is small, then $\hat{X}_i^{(c)}(t-1)$ will be positive with high probability. A result analogous to Proposition 2 holds if instead $\hat{X}_i^{(c)}(t-1) < 0$ for $c > 1$. It is conjectured but not proven in this paper that $q_i(t)$ is monotonically decreasing in $\lambda_i(t-1)$ under much weaker conditions than stated in Proposition 2.

Bounds on the objective function proposed in this section will be given in the next section.

III. Performance bounds

This section presents bounds on the performance gain that can be achieved using adaptive resource allocation policies. These bounds are used to quantify the benefit of adaptive sensing, particularly in the presence of targets with varying mission importance. Lower bounds on the objective function are obtained

by analyzing the performance of an oracle policy, which are optimal allocations that depend on exact knowledge of the locations and (in some cases) the classes of the targets. An upper bound is obtained by analyzing the performance of a baseline non-adaptive policy that allocates resources uniformly over the scene, called the uniform policy.

Assumption 1 on the equality of prior variances σ_c^2 is used in this section in order to derive bounds in Propositions 3–5 that are closed-form and facilitate interpretation. Exact expressions that do not require Assumption 1 are provided in Appendix X. A comparison of the two approaches in Fig. 3 shows that the main conclusions of this section remain valid without Assumption 1.

In all of the policies considered in this section, $\boldsymbol{\lambda}$ is a deterministic function of the target classes $\boldsymbol{C}$. For example, in the oracle policy, the target class is known, and the same allocation is given to targets of the same class, while in the uniform policy, the same allocation is given to targets of all classes. In this case, the objective function can be further simplified, as given below:

Lemma 1 When $\boldsymbol{\lambda}$ is a deterministic (non-random) function of $\boldsymbol{C}$,

$$J_T(\boldsymbol{\lambda}) = \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\sum_{i=1}^{N} \frac{h(C_i)}{\nu^2/\sigma_c^2 + \overline{\lambda_i}} \right]. \tag{18}$$

Proof: The proof is given in Appendix V. ■

A. Full oracle policy

The full oracle policy has exact knowledge of both the locations and classes of targets within the scene. In other words, $p_i^{(c)}(t) = I_i^{(c)}$ for all $i = 1, 2, \ldots, N$, $c \in \mathcal{C}$, and $t = 0, 1, \ldots, T-1$. The optimal full-oracle policy is given by

$$\overline{\lambda^o_{\pi(i)}} = \begin{cases} \dfrac{\left(\Lambda + \nu^2 \sum\limits_{j=1}^{k^*} \sigma^{-2}_{C_{\pi(j)}}\right) \sqrt{h(C_{\pi(i)})}}{\sum_{j=1}^{k^*} \sqrt{h(C_{\pi(j)})}} - \dfrac{\nu^2}{\sigma^2_{C_{\pi(i)}}}, & i = 1, \ldots, k^* \\ 0, & \text{otherwise} \end{cases}. \tag{19}$$

where the number of non-zero allocations k^* depends on SNR, and π is an index permutation that sorts $\sqrt{h(C_i)}\sigma^2_{C_i}$ in non-increasing order so that

$$\sqrt{h(C_{\pi(1)})}\sigma^2_{C_{\pi(1)}} \geq \sqrt{h(C_{\pi(2)})}\sigma^2_{C_{\pi(2)}} \geq \cdots \geq \sqrt{h(C_{\pi(N)})}\sigma^2_{C_{\pi(N)}}. \tag{20}$$

The derivation of the optimal full-oracle policy is given in Appendix VI.

To analyze the performance of the full oracle policy, we make the following definitions and assumptions. Denote the number of targets with class c as

$$N_c = |\{i : C_i = c\}| \tag{21}$$

Define the first and second moments of $\sqrt{h(C_i)}$ conditioned on $C_i > 1$ (so that $h(C_i) > 0$):

$$m_{h,1} \stackrel{\triangle}{=} \mathbb{E}_{C_i>1}\left[\sqrt{h(C_i)}\right] = \frac{1}{\bar{p}_1}\sum_{c=2}^{|\mathcal{C}|} p_c\sqrt{h(c)} \tag{22}$$

$$m_{h,2} \stackrel{\triangle}{=} \mathbb{E}_{C_i>1}\left[h(C_i)\right] = \frac{1}{\bar{p}_1}\sum_{c=2}^{|\mathcal{C}|} p_c h(c) \tag{23}$$

where $\bar{p}_1 = 1 - p_1$. We invoke Assumption 1 and the following:

Assumption 2 The total resource budget is greater than $\Lambda_{\min}$,

$$\Lambda_{\min} \geq c_0 \max\{\Lambda_0, \Lambda_1\}, \tag{24}$$

where $c_0 = \nu^2/\sigma_0^2$ and

$$\Lambda_0 = \sum_{c=2}^{|\mathcal{C}|} N_c\sqrt{h(c)/h(2)} - (N - N_1) \tag{25}$$

$$\Lambda_1 = \left(m_{h,2}/m_{h,1}^2\right) - 1 \tag{26}$$

The condition $\Lambda_{\min} \geq c_0\Lambda_0$ ensures that at least some effort is distributed to all classes with non-zero mission importance (i.e. $h(c) > 0$). This allows (19) to be rewritten as

$$\overline{\lambda_i^o} = \frac{(\Lambda + (N - N_1)c_0)\sqrt{h(C_i)}}{\sum_{c\in\mathcal{C}} N_c\sqrt{h(c)}} - c_0, \tag{27}$$

when $C_i > 1$ and 0 otherwise. The second condition $\Lambda_{\min} \geq c_0\Lambda_1$ implies a certain functional convexity that facilitates bounding expectations. Note that Λ_0 in (25) depends on the histogram of classes $\boldsymbol{N} = \{N_c\}_{c=1,2,\ldots,|C|}$, which is random. In Appendix VII, we show that if Λ_0 in (24) is replaced by the deterministic quantity

$$\Lambda_0' = \alpha N\bar{p}_1\left(\frac{m_{h,1}}{\sqrt{h(2)}} - 1\right), \tag{28}$$

where $\alpha > 1$ is an arbitrary constant, then the original condition $\Lambda_{\min} \geq c_0\Lambda_0$ is satisfied with probability converging exponentially to 1 as N increases.

Proposition 3 Given Assumptions 1 and 2, the cost (18) of the full oracle policy (19) is bounded from below $J_T(\boldsymbol{\lambda}^o) \geq L_T(\boldsymbol{\lambda}^o)$ with

$$L_T(\boldsymbol{\lambda}^o) \stackrel{\triangle}{=} \frac{\nu^2 N\bar{p}_1}{\Lambda + N\bar{p}_1 c_0}(m_{h,2} + (N\bar{p}_1 - 1)m_{h,1}^2), \tag{29}$$

and bounded from above $J_T(\boldsymbol{\lambda}^o) \leq U_T(\boldsymbol{\lambda}^o)$ with

$$\begin{aligned} U_T(\boldsymbol{\lambda}^o) &\stackrel{\triangle}{=} L_T(\boldsymbol{\lambda}^o) + \nu^2 \frac{((\Lambda + c_0)m_{h,1}^2 - c_0 m_{h,2})Np_1\bar{p}_1}{\Lambda^2} \\ &= \frac{\nu^2 N\bar{p}_1}{\Lambda + N\bar{p}_1 c_0}\left(m_{h,2} + (N\bar{p}_1 - 1 + O(1))m_{h,1}^2\right) \end{aligned} \tag{30}$$

Proof: The proof is given in Appendix II. ■

B. Location-only oracle policy

A location-only oracle policy has knowledge of the locations of targets with $h(C_i) > 0$, but not the classes themselves. The location-only oracle policy is

$$\overline{\lambda_i^{lo}} = \begin{cases} \Lambda/(N - N_1), & h(C_i) > 0 \\ 0, & \text{otherwise} \end{cases} \tag{31}$$

This policy was studied in previous work [3,14] for the case $|\mathcal{C}| = 2$. The proposition below gives a lower bound (32) on the cost of the policy that generalizes a lower bound in [14, Sec. III-B] for the 2-class case, and additionally provides an upper bound (33) that converges to the lower bound for large N.

Proposition 4 Given Assumption 1, the cost of the location-only oracle policy is bounded from below with $J_T(\boldsymbol{\lambda}^{lo}) \geq L_T(\boldsymbol{\lambda}^{lo})$ with

$$L_T(\boldsymbol{\lambda}^{lo}) \stackrel{\triangle}{=} \frac{\nu^2 N^2 \bar{p}_1^2 m_{h,2}}{\Lambda + N\bar{p}_1 c_0}. \tag{32}$$

and bounded from above with $J_T(\boldsymbol{\lambda}^{lo}) \leq U_T(\boldsymbol{\lambda}^{lo})$ with

$$\begin{aligned} U_T(\boldsymbol{\lambda}^{lo}) &\stackrel{\triangle}{=} L_T(\boldsymbol{\lambda}^{lo}) + \frac{Np_1\bar{p}_1}{\Lambda}(\nu^2 m_{h,2}) \\ &= \frac{\nu^2 N\bar{p}_1}{\Lambda + N\bar{p}_1 c_0}\left(N\bar{p}_1 + O(1)\right) m_{h,2} \end{aligned} \tag{33}$$

Proof: The proof is given in Appendix VIII. ■

C. Uniform (baseline) policy

The uniform policy uniformly allocates resources everywhere in the scene, :

$$\lambda_i^u(t) = \Lambda(t)/N \tag{34}$$

, and the allocations across all stages are

$$\overline{\lambda_i^u} = \sum_{t=0}^{T-1} \lambda_i^u(t) = \Lambda/N \tag{35}$$

Proposition 5 Given Assumption 1, the cost of the uniform policy is

$$J_T(\boldsymbol{\lambda}^u) = \frac{\nu^2 N \bar{p}_1 m_{h,2}}{c_0 + \Lambda/N} \tag{36}$$

Proof: The proof is given in Appendix IX. ■

D. Comparisons of bounds

The bounds in Propositions 3–5 are compared to yield insight into the regimes where one might expect significant gains in comparison to the non-adaptive baseline (uniform) policy $\boldsymbol{\lambda}^u$, as well as the regimes in which incorporating mission importance (through the full oracle policy $\boldsymbol{\lambda}^o$) yields significant gains over a detection-only policy (i.e., the location-only oracle $\boldsymbol{\lambda}^{lo}$). The model parameters of interest are the total effort budget Λ or equivalently SNR[1], the prior distribution of targets $\{p_c\}_{c\in\mathcal{C}}$, and the mission importance $\{h(c)\}_{c\in\mathcal{C}}$. Note that $\{p_c, h(c)\}_{c\in\mathcal{C}}$ determine the moments $m_{h,1}$ and $m_{h,2}$ on which the bounds depend directly.

Define $G(\boldsymbol{\lambda}^0, \boldsymbol{\lambda}^1)$ as the gain of policy $\boldsymbol{\lambda}^1$ with respect to another policy $\boldsymbol{\lambda}^0$:

$$G(\boldsymbol{\lambda}^0, \boldsymbol{\lambda}^1) = J_T(\boldsymbol{\lambda}^0)/J_T(\boldsymbol{\lambda}^1). \tag{37}$$

From Propositions 3–5, we obtain the following bounds on the performance gains of the oracle policies with respect to the uniform policy:

$$G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^o) \leq \left(\frac{\Lambda + N\bar{p}_1 c_0}{N c_0 + \Lambda}\right) g_1(N, \bar{p}_1, m_{h,1}, m_{h,2}) \tag{38}$$

$$G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{lo}) \leq \left(\frac{\Lambda + N\bar{p}_1 c_0}{N c_0 + \Lambda}\right) g_2(\bar{p}_1) \tag{39}$$

[1]SNR is defined as a mapping from the total budget Λ: $\Lambda = 10^{\mathrm{SNR}/10} N$.

where the functions g_1 and g_2 are defined as

$$g_1(N, \bar{p}_1, m_{h,1}, m_{h,2}) = \frac{N m_{h,2}}{m_{h,2} + (N\bar{p}_1 - 1) m_{h,1}^2} \tag{40}$$

$$g_2(\bar{p}_1) = (\bar{p}_1)^{-1} \tag{41}$$

The location-only oracle bound (39) coincides with the oracle bound [14, eq. (16)] in the 2-class case.

The first term in both (38) and (39) approaches 1 as the SNR $\Lambda \gg 1$. In this regime, the gain of the location-only oracle depends only on the proportion of non-zero targets through $\bar{p}_1$. To further interpret the full oracle bound (38), we observe that $m_{h,1}$ is the first moment of the random variable $\sqrt{h(C_i)}$ conditioned on $C_i > 1$, while $m_{h,2}$ is the corresponding second moment. Then, since the second moment is greater than the first moment squared, $m_{h,1}^2 \leq m_{h,2}$, we have

$$\begin{aligned} g_1(N, \bar{p}_1, m_{h,1}, m_{h,2}) &\leq \frac{N m_{h,2}}{m_{h,1}^2 + (N\bar{p}_1 - 1) m_{h,1}^2} \\ &= \frac{m_{h,2}}{\bar{p}_1 m_{h,1}^2}, \end{aligned} \tag{42}$$

where equality holds asymptotically as $N \to \infty$. The ratio $m_{h,2}/m_{h,1}^2$ can be interpreted as $1 + \mathrm{CV}_{C_i>1}(\sqrt{h(C_i)})$, where the latter is the coefficient of variation (CV, a normalized measure of variability) of the mission importance $\sqrt{h(C_i)}$ conditioned on $C_i > 1$.

Figure 1 evaluates the bound (38) on the gain of the full oracle policy with respect to the uniform policy $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^o)$ across the model parameters. Unless otherwise stated, the simulation parameters are given in Table I. It is seen that decreasing $\bar{p}_1$ (leading to fewer overall targets) and increasing SNR, N and the ratio $(m_{h,2}/m_{h,1}^2)$ all lead to larger gains. Moreover, the gains incurred by increasing N and/or SNR approach a finite value (i.e., diminishing returns given sufficiently high N or SNR), while they continue to grow with decreasing $\bar{p}_1$ or increasing $(m_{h,2}/m_{h,1}^2)$.

Next, the bounds on the full and location-only oracles are compared directly to obtain

$$\begin{aligned} G(\boldsymbol{\lambda}^{lo}, \boldsymbol{\lambda}^o) &\leq \frac{(N\bar{p}_1 + O(1)) m_{h,2}}{m_{h,2} + (N\bar{p}_1 - 1) m_{h,1}^2} \\ &\overset{N\to\infty}{\to} \frac{m_{h,2}}{m_{h,1}^2} = 1 + \mathrm{CV}_{C_i>1}(\sqrt{h(C_i)}). \end{aligned} \tag{43}$$

Similarly, we have

$$\begin{aligned} G(\boldsymbol{\lambda}^{lo}, \boldsymbol{\lambda}^o) &\geq \frac{(N\bar{p}_1) m_{h,2}}{m_{h,2} + (N\bar{p}_1 - 1 + O(1)) m_{h,1}^2} \\ &\overset{N\to\infty}{\to} \frac{m_{h,2}}{m_{h,1}^2} = 1 + \mathrm{CV}_{C_i>1}(\sqrt{h(C_i)}). \end{aligned} \tag{44}$$

TABLE I

SIMULATION PARAMETERS

Parameter	Value
Number of locations, N	2500
Number of classes, $\lvert\mathcal{C}\rvert$	3
Class probabilities, $\{p_c\}_{c\in\mathcal{C}}$	$\{0.95, 0.049, 0.001\}$
Class importance, $\{h(c)\}_{c\in\mathcal{C}}$	$\{0, 1, 2500\}$
Target prior means, $\{\mu_c\}_{c\in\mathcal{C}}$	$\{0, 3, 1.5\}$
Target prior variances, $\{\sigma_c^2\}_{c\in\mathcal{C}}$	$\{0, 1/16, 1/16\}$
Noise variance, ν^2	1
Number of sensors for LA policy, M	400
Number of sensors for GU/LA policy, M	50

Thus, when N is large, the upper and lower bounds converge to the same value (i.e., 1 plus the CV). The CV is equal to zero in the 2-class case or when the target classes have the same importance. In these cases, the full and location-only oracles coincide. Otherwise, the CV is positive and the asymptotic gain is greater than 1, indicating that the full-oracle performs better than the location-only oracle. Hence the CV can be seen as representing the value of knowledge of target importance.

Figure 2 evaluates the bound (43) on the gain of the full oracle policy with respect to the location-only oracle policy $G(\boldsymbol{\lambda}^{lo}, \boldsymbol{\lambda}^{o})$ across model parameters. Increasing the ratio $(m_{h,2}/m_{h,1}^2)$ leads to larger gains, as expected from (43), (44) and similar to Fig. 1(b). Increasing N also leads to marginally larger gains. However, increasing $\bar{p}_1$ (i.e., increasing the number of targets) increases the gains between the oracle policies, which is the opposite pattern as compared to Fig. 1(b). This suggests that when there is either a large spread in mission importance or many targets of interest, the full oracle policy has large gains over the location-only oracle.

Finally we numerically evaluate the impact of Assumption 1 on the analysis of the full oracle policy. Figure 3 shows the ratio of the analytical upper bound (30) on the full oracle cost to the exact expression (18) without Assumption 1. The ratio is always greater than one (above 0 dB) since Assumption 1 gives an upper bound on the actual cost. It is seen that when (a) σ_2^2 is either 100 times larger or smaller than σ_3^2, and (b) SNR is low, there are gaps but the ratio never exceeds 2 (3 dB).

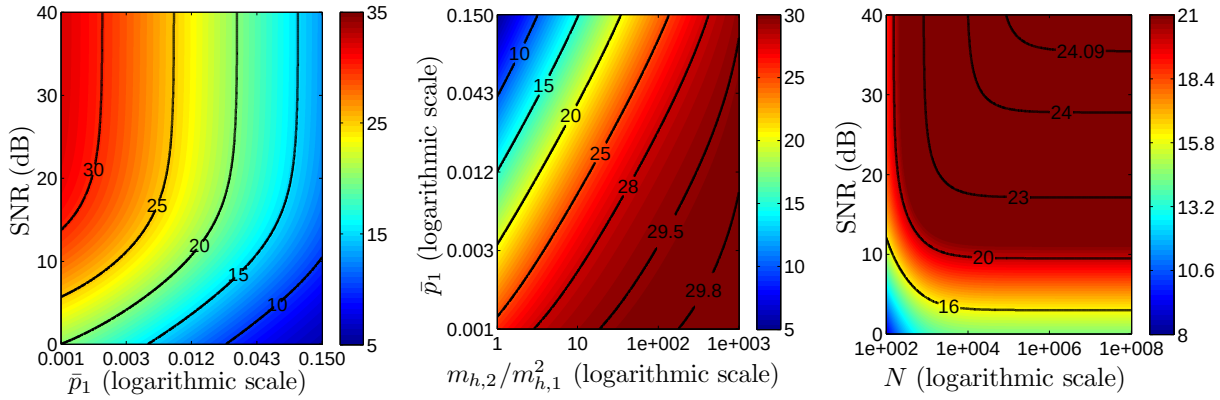

(a) (38) vs. ($\bar{p}_1$,SNR) (b) (38) vs. $\left(\frac{m_{h,2}}{m_{h,1}^2}, \bar{p}_1\right)$ (c) (38) vs. (N,SNR)

Fig. 1. Bound (38) on the gain $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^o)$ of the full oracle policy w.r.t. the uniform policy across the model parameters. The bound is plotted in units of dB ($10\log_{10}$). Decreasing $\bar{p}_1$ (leading to fewer overall targets) and increasing SNR, N and the ratio $(m_{h,2}/m_{h,1}^2)$ lead to larger gains.

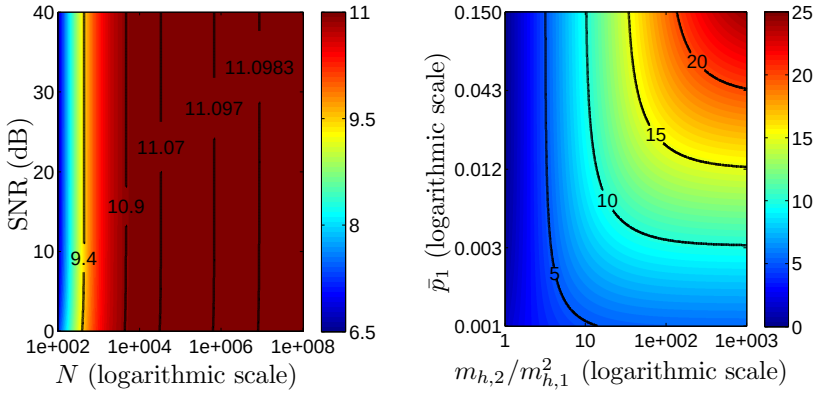

(a) (43) vs. (N,SNR) (b) (43) vs. $\left(\frac{m_{h,2}}{m_{h,1}^2}, \bar{p}_1\right)$

Fig. 2. Bound ((43), in dB) on the gain $G(\boldsymbol{\lambda}^{lo}, \boldsymbol{\lambda}^o)$ of the full oracle policy with respect to the location-only oracle policy. Increasing the ratio $(m_{h,2}/m_{h,1}^2)$ leads to larger gains as in Fig. 1(b). Increasing N also leads to marginally larger gains. However, increasing $\bar{p}_1$ (i.e., increasing the number of targets) increases the gains between the oracle policies, which is the opposite pattern as compared to Fig. 1(b).

IV. Search Policies

This section develops implementable policies that adaptively allocate resources to locations with targets. Recall a policy $\{\boldsymbol{\lambda}(t)\}_{t=0}^{T-1}$ is a sequence of effort allocations, which are mappings from previous observations $\mathbf{Y}(t)$ to $[0, \Lambda(t)]^N$. Denote $\mathcal{U}(t)$ as the set of locations i for which $\lambda_i(t) > 0$. For example, $\mathcal{U}(t) = \{1, \cdots, N\}$ for the uniform policy and $\mathcal{U}(t) = \{i | C_i > 1\}$ for the full oracle policy.

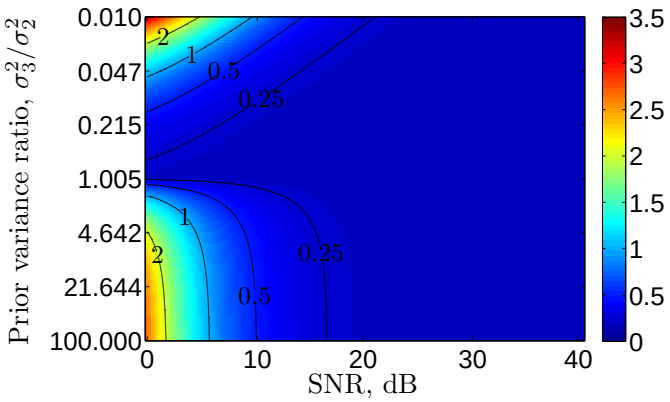

Fig. 3. Ratio (in dB) of the upper bound (30) on the full oracle cost to its exact value computed without Assumption 1. There is no significant difference when σ_2^2 is within a factor of 100 of σ_3^2 or SNR is moderate to high.

We introduce several models to emulate realistic types of sensors. Sensors typically vary in sensing agility, e.g. field of view and accuracy, and in cost. Specifically, three types of sensors are considered: the global uniform (GU) sensor, the global adaptive (GA) sensor, and the local adaptive (LA) sensor. The global adaptive sensor can choose any subset of the scene and allocate different amounts of sensing resource to the selected locations. These sensors provide the highest agility, but may be difficult or very expensive to realize. The global uniform sensor explores all locations with equal sensing resource in each location; e.g., a large field of view radar. Therefore, it is likely to spend resources in locations without targets and suffer performance degradation compared to the global adaptive sensors. Finally, the local adaptive sensor can only explore a small number of locations within the scene, albeit with high resolution, e.g., unmanned aerial vehicles (UAVs). These sensors may be easily available in practice and are simpler to deploy as compared to global adaptive sensors. Table II compares these sensor types with regard to their allocations, ease of implementation, and ability to adapt.

A. Global adaptive (GA) policy

As discussed in previous work on similar problems [8], it is possible in principle to use dynamic programming (DP) to globally optimize $J_T(\boldsymbol{\lambda})$ with respect to the entire allocation policy $\boldsymbol{\lambda}$ subject to the resource constraint (4). However, for $T > 2$, this exact solution is computationally intractable. As an alternative we present a myopic method that determines each $\boldsymbol{\lambda}(t)$ independently for $t = 0, 1, \ldots, T-1$

TABLE II

COMPARISON OF SENSOR TYPES AND ALLOCATIONS

Sensor Type	Allocation	Adaptivity	Implementation
Global Adaptive	$\lambda_i(t) \in [0, \Lambda/T]$, $\mathcal{U}(t)$ any subset	Full	Difficult
Global Uniform	$\lambda_i(t) = \Lambda/(NT)$, $\mathcal{U}(t) = \{1, \ldots, N\}$	None	Easy
Local Adaptive	$\lambda_i(t) = k\Lambda/(MT)$, $k = \{0, 1, \ldots, M\}$, $\lvert\mathcal{U}(t)\rvert \leq M$	Limited	Medium

as the solution to

$$\boldsymbol{\lambda}^{ga}(t) = \underset{\boldsymbol{\lambda}(t)}{\operatorname{argmin}} \; K(t; \boldsymbol{\lambda}(t)) \quad \text{s.t.} \; \sum_{i=1}^{N} \lambda_i(t) = \Lambda(t), \tag{45}$$

where the myopic cost at stage t given $\boldsymbol{Y}(t)$ is defined as

$$K(t; \boldsymbol{\lambda}(t)) = \sum_{i=1}^{N} \sum_{c=2}^{|\mathcal{C}|} \frac{p_i^{(c)}(t) h(c)}{\nu^2/(\sigma_i^{(c)}(t))^2 + \lambda_i(t)}, \tag{46}$$

and $\Lambda(t)$ is some fraction of the total budget Λ. In this work, Λ is equally allocated across stages, $\Lambda(t) = \Lambda/T$. The myopic cost (46) is inspired by the exact cost function (15) and is equivalent to (15) when $t = T - 1$. However, (46) requires no expectation because we condition on (i.e. have collected and incorporated) measurements up to stage t.

The myopic allocation problem (45) is a convex optimization because the cost function is a sum of convex functions of the form $a/(b + \lambda_i(t))$, $a, b \geq 0$, and the constraints are linear (recall that $\lambda_i(t)$ is also constrained to be non-negative). Hence it may be solved efficiently using a variety of iterative algorithms. Furthermore, the optimal solution $\boldsymbol{\lambda}^{ga}(t)$ satisfies the following precedence property: if

$$\sum_{c=2}^{|\mathcal{C}|} p_i^{(c)}(t) h(c) \big(\sigma_i^{(c)}(t)\big)^4 \leq \sum_{c=2}^{|\mathcal{C}|} p_j^{(c)}(t) h(c) \big(\sigma_j^{(c)}(t)\big)^4, \tag{47}$$

then we cannot have $\lambda_i^{ga}(t) > 0$ and $\lambda_j^{ga}(t) = 0$. This property may be derived in a manner similar to [8, App. C] by invoking the optimality condition for (45). The property implies that if $\boldsymbol{\lambda}^{ga}(t)$ has k nonzero components, they must correspond to the k largest of the quantities appearing in (47).

Algorithm 1 $\{\boldsymbol{\lambda}^{ga}(t-1), \boldsymbol{\xi}(t)\}_{t=1}^{T}$ =GlobalAdapt$(\Lambda, \boldsymbol{\xi}(0); T)$

0. Set $\Lambda(t) = \Lambda/T$ for $t = 0, 1, \ldots, T-1$.

for $t = 1, 2, \ldots, T$ **do**

1. Calculate $\lambda_i^{ga}(t-1)$ through convex optimization (45).

2. Take measurements $y_i(t)$ as in (3).

3. Update posterior dist. $\boldsymbol{\xi}(t)$ with (9),(10),(11).

end for

Under Assumption 1, the posterior variances do not depend on class, $(\sigma_i^{(c)}(t))^2 = \sigma_i^2(t)$. In this case, by defining

$$z_i(t) = \sum_{c=2}^{|\mathcal{C}|} p_i^{(c)}(t) h(c), \tag{48}$$

the optimal solution $\boldsymbol{\lambda}^{ga}(t)$ can be expressed analytically as in [8]. We refer the reader to eqs. (21)–(24) of [8] with $p_i(t)$ replaced by $z_i(t)$.

The GA policy including resource allocation, measurement, and posterior update steps is summarized in Algorithm 1 with inputs Λ and prior belief state $\boldsymbol{\xi}(0)$, where $\boldsymbol{\xi}(t)$ is defined as

$$\boldsymbol{\xi}(t) = \left\{ p_i^{(c)}(t), \hat{X}_i^{(c)}(t), (\sigma_i^{(c)}(t))^2 \right\}_{i=1,\ldots,N, c\in\mathcal{C}} \tag{49}$$

B. Local adaptive (LA) policy

In some cases it may be impossible to deploy a sensor with the agility to assign different resources to every location in the scene at every stage t. Instead we consider the situation where there are M local sensors (e.g. UAVs) that can explore a subset of the locations with a fixed resource amount per sensor. At each stage $t = 0, 1, \ldots, T-1$, define $u_i(t) \in \{0, 1, \ldots, M\}$ as the number of sensors allocated to location i for $i \in \mathcal{X}$, $\sum_{i=1}^{N} u_i(t) = M$, and $\boldsymbol{u}(t) = \{u_i(t)\}_{i=1}^{N}$. Define the local adaptive allocation given $\boldsymbol{u}(t)$ as

$$\lambda_i^{la}(t; \boldsymbol{u}(t)) = u_i(t) \left(\frac{\Lambda(t)}{M} \right) \tag{50}$$

Then we seek the allocation that minimizes (46):

$$\hat{\boldsymbol{u}}(t) = \underset{\boldsymbol{u}(t)}{\operatorname{argmin}} \; K(t; \boldsymbol{\lambda}^{la}(t; \boldsymbol{u}(t))) \tag{51}$$

for $t = 0, 1, 2, \ldots, T-1$.

Algorithm 2 $\left\{\boldsymbol{\lambda}^{la}(t-1), \boldsymbol{\xi}(t)\right\}_{t=1}^{T}$ =LocalAdapt($\Lambda, \boldsymbol{\xi}(0); T$)

0. Set $\Lambda(t) = \Lambda/T$ for $t = 0, 1, \ldots, T-1$.

for $t = 1, 2, \ldots, T$ **do**

1. Set $\hat{u}_i(t-1) = 0$ for $i = 1, 2, \ldots, N$.

2. Calculate $\Delta_i(t-1)$ for $i = 1, \ldots, N$ using (52).

for $m = 1, 2, \ldots, M$ **do**

3. Choose $i^* = \operatorname{argmax}_i \Delta_i(t-1)$.

4. Set $\hat{u}_{i^*}(t-1) \leftarrow \hat{u}_{i^*}(t-1) + 1$.

5. Re-calculate $\Delta_{i^*}(t-1)$ using (52).

end for

6. Calculate $\lambda_i^{la}(t-1; \boldsymbol{u}(t-1))$ with (50).

7. Take measurements $y_i(t)$ as in (3).

8. Update posterior dist. $\boldsymbol{\xi}(t)$ with (9), (10), (11).

end for

The solution is given by a series of M steps, where at each step, we allocate a single sensor to the location that provides the greatest decrease in (46). First, set $\hat{u}_i(t) = 0$ for $i = 1, 2, \ldots, N$. Then, define the decrease in cost by allocating a single additional sensor to location i as

$$\begin{aligned}\Delta_i(t) = &\sum_{c=2}^{|\mathcal{C}|} \frac{p_i^{(c)}(t) h(c)}{\nu^2/(\sigma_i^{(c)}(t))^2 + \hat{u}_i(t)\Lambda(t)/M} \\ &- \sum_{c=2}^{|\mathcal{C}|} \frac{p_i^{(c)}(t) h(c)}{\nu^2/(\sigma_i^{(c)}(t))^2 + (\hat{u}_i(t)+1)\Lambda(t)/M}.\end{aligned} \tag{52}$$

In each of M steps, we set $\hat{u}_{i^*}(t) \leftarrow \hat{u}_{i^*}(t) + 1$, where i^* is the index with the largest $\Delta_i(t)$. Note that multiple sensors are allowed to visit the same location.

The greedy allocation above is optimal for (51) because of the convexity of the myopic cost with respect to each $\lambda_i(t)$. As a consequence, the decreases $\Delta_i(t)$ at a particular location diminish as more sensors are assigned to it, and we can be sure that the assignment in each step yields the largest decrease globally.

The LA policy is summarized in Algorithm 2.

C. Global uniform/local adaptive (GU/LA) mixture policy

With no prior information on the location of targets in the scene, the local policy is likely to perform poorly, because it may take a long time to locate targets if only a small number of positions can be queried in each stage. Thus, we consider a third sensing modality where a global sensor is able to gather low resolution information on target location by uniformly observing the scene. Subsequently, local sensors can use this information to measure the likely locations with higher signal quality. In this policy, the optimization is two-fold: (a) the percentage of resources used by the global exploration sensor is optimized; and (b), the locations of the local sensors are optimized. When $|\mathcal{U}(t)| = M$ for all t, this optimization problem can be written as

$$\boldsymbol{\lambda}^{gula} = \underset{T_s, \{\boldsymbol{u}(t)\}_{t=0}^{T-1}}{\operatorname{argmin}} J_T(\boldsymbol{\lambda}^{gula}(\{\boldsymbol{u}(t)\}_{t=0}^{T-1}, T_s)) \tag{53}$$

$$\text{where} \quad \lambda_i^{gula}(t; \boldsymbol{u}(t), T_s) = \begin{cases} \Lambda(t)/N, & t \leq T_s \\ u_i(t)\Lambda(t)/M, & t > T_s \end{cases}$$

Given T_s, the optimization problem reduces to finding the sensor allocations $\boldsymbol{u}(t)$ for $t = T_s + 1, \ldots, T-1$, which is done myopically using steps (1-5) from the Local Adaptive policy (Algorithm 2). Optimization over T_s is done offline (i.e., before any real measurements are taken) by approximating the expectation in $J_T(\boldsymbol{\lambda})$ through Monte Carlo samples. In general, this requires $\mathcal{O}(T)$ Monte Carlo samples to determine the optimal T_s in a T-stage policy. The GU/LA algorithm is summarized by Algorithms 3 and 4.

V. Simulation

This section presents a numerical study for comparing the performance of proposed policies — global adaptive (GA), local adaptive (LA), and global uniform/local adaptive mixture (GU/LA) — and a previously proposed policy ARAP [3,8], which is designed to only detect targets and not classify them. The global uniform (GU) and full oracle policies discussed in Section III are used as benchmarks.

Simulation parameters are given in Table I unless otherwise stated. In this scenario, there are very few targets in the scene (5% on average), which leads to significant performance gaps between the GU policy and adaptive policies. Moreover, the condition that $\mu_3 < \mu_2$ is imposed to account for the fact that high-importance targets are generally harder to detect than low-importance targets. We set $T = 10$ for the GA policy and $T = 30$ for the GU/LA and LA policies since the latter two require additional stages to effectively search the entire scene. Note that once the SNR is fixed, the total budget Λ is the same

Algorithm 3 $\left\{\boldsymbol{\lambda}^{gula}(t-1), \boldsymbol{\xi}(t)\right\}_{t=1}^{T}$ =GU/LA$(\Lambda, \boldsymbol{\xi}(0); T)$

0. Set $\Lambda(t) = \Lambda/T$ for $t = 0, 1, \ldots, T-1$.

1. Find optimal T_s using Algorithm 4

for $t = 1, 2, \ldots, T$ **do**

if $t \leq T_s$ **then**

2-a. $\lambda_i(t-1) = \Lambda(t)/N$ for $i = 1, 2, \ldots, N$

else

2-b. $\lambda_i(t-1)$ given by steps (1)-(6) in Alg. 2.

end if

3. Take measurements $y_i(t)$ as in (3)

4. Update posterior dist. $\boldsymbol{\xi}(t)$ with (9), (10), (11)

end for

regardless of T. One could additionally optimize over the per-stage budget allocations, $\{\Lambda(t)\}_{t=0,1,\ldots,T-1}$ as in [8], though this is not done in the current paper.

This section begins by identifying the regimes in which the GA policy nearly achieves the oracle policy performance as a function of the model parameters, as well as the regimes where there are few performance gains over the baseline GU policy. The performance of the GU/LA policy is analyzed over the same parameters. Additionally, this section explores the sensitivity of the GU/LA policy to (a) correct specification of the percentage of stages used for global sensing, T_s/T, and (b) the number of local sensors available. Subsequently, we demonstrate the performance improvement by including the global sensor by analyzing the differences between the GU/LA and LA policies. The section concludes by comparing the performance of the proposed policies across other performance metrics which are not directly optimized, including mean-squared error and misclassification probability.

In Fig. 4, the benefits of adaptive sensing are explored by comparing the gain (37) of the GA, GU/LA, and oracle policies to the baseline GU policy. Figs. 4a, 4c, and 4e show the gains when varying SNR and the high-value target class probability p_3 and fixing $h(3) = 2500$. Meanwhile, Figs. 4b, 4d, and 4f show the gains when varying p_3 and the high-value target mission importance $h(3)$ and fixing SNR to 20 dB. Black dots indicate that the gains of the GA and GU/LA policies are within 3 dB of the oracle gain. The GA and GU/LA policies have improved gains when either the SNR increases, the mission importance increases, or the number of high-value targets decreases. At low SNR levels (about 0 dB), the GA and

Algorithm 4 T_s =findOptimalGULAPoint$(\Lambda, \boldsymbol{\xi}(0); T)$

0. Set $\Lambda(t) = \Lambda/T$ for $t = 0, 1, \ldots, T-1$.

for $T_s \in \{1, \ldots, T\}$ **do**

for $j = 1, 2, \ldots, N_{\mathrm{MC}}$ **do**

1. Draw Monte Carlo sample of $\{C_i, X_i\}_{i=1}^{N}$ from $\boldsymbol{\xi}(0)$.

2. Set $\boldsymbol{\xi}^{(j)}(0) = \boldsymbol{\xi}(0)$.

for $t = 1, 2, \ldots, T$ **do**

if $t \leq T_s$ **then**

3-a. $\lambda_i(t-1) = \Lambda(t)/N$ for $i = 1, 2, \ldots, N$

else

3-b. $\lambda_i(t-1)$ given by steps (1)-(6) in Alg. 2.

end if

4. Draw $\boldsymbol{y}(t)$ from (3) given $\boldsymbol{C}$, $\boldsymbol{X}$, and $\boldsymbol{\lambda}(t-1)$

5. Update posterior dist. $\boldsymbol{\xi}^{(j)}(t)$ with (9)-(11)

end for

6. Calculate cost for (j)-th sample, $Q(j, T_s) = K(T; \boldsymbol{\lambda}(T) = \mathbf{0})$ using (46).

end for

7. Calculate $J_T(T_s) \approx (1/N_{\mathrm{MC}}) \sum_{j=1}^{N_{MC}} Q(j, T_s)$.

end for

8. Return $T_s = \mathrm{argmin}_{T_s'} J_T(T_s')$.

GU/LA policies have few gains over the GU policy. However, for sufficiently high SNR levels (≥ 15 dB for the GA policy, and ≥ 20 dB for the GU/LA policy), the policies nearly achieve the performance of the oracle policy over all other model parameters. Note that the GA and GU/LA policies always perform better than the GU policy (i.e., gains are always positive).

Figs. 5a and 5b explore the sensitivity of the GU/LA policy to the percentage of total resources used by the global sensor, T_s/T. In Fig. 5a, it is seen that there is significant performance degradation when this percentage is close to either extreme. When $T_s = T$ (i.e., the GU policy), the lack of adaptivity leads to inefficient resource allocation. Conversely, when $T_s = 0$ (i.e., the LA policy), the local sensors have difficulty in finding the locations that contain valuable targets. Nevertheless, for any fixed SNR, the gain of the GU/LA policy is relatively flat in a region around the optimal T_s value. Circles indicate

the maximum T_s/T within 3 dB of the maximum gain, while diamonds indicate the minimum T_s/T within 3 dB. It is seen that for any SNR, there is a large region where the gain is within 3 dB, which suggests that the GU/LA policy may be robust to small errors in finding the optimal T_s. Fig. 5b shows the optimal T_s values while varying both SNR and p_3. It is seen that most of the variation occurs when SNR changes. Indeed, this was also true when fixing SNR and varying p_3 and $h(3)$ (not shown), where it was found that the optimal T_s was nearly constant over all values.

Fig. 6 explores the benefit of including a global sensor by comparing the GU/LA and LA policies. Figs. 6a and 6b show the gains of the LA and GU/LA policies (37) with respect to the GU policy, while varying SNR and the number of local sensors. In this simulation, the optimal T_s for the GU/LA policy is given by the values in Fig. 5a. Black diamonds indicate the minimum number of sensors needed to achieve within 3 dB of the maximum gain at each SNR. Note that the maximum number of sensors is equal to $N = 2500$ for both policies. However, the LA policy requires at least 100 local sensors in almost all cases to attain performance within 3 dB of the maximum, while the GU/LA policy requires an order of magnitude fewer sensors. Furthermore, a phase transition occurs for the LA policy when fewer than 100 sensors are used, wherein the LA policy actually performs worse than the GU policy.

Figs. 6c and 6d show the performance of the LA and GU/LA policies for various number of sensors as compared to the GA and oracle policies. Similar to above, the LA policy requires many more sensors than the GU/LA policy in order to achieve performance close to the GA policy.

The proposed policies approximately optimize an objective function that combines estimation and classification errors. However, the policies also perform well in reducing each of these errors individually as demonstrated in the next figures. Fig. 7 plots the posterior variance (i.e., expected estimation error) within the low-value and high-value target classes. Fig. 8 shows the misclassification probabilities within each class. Fig. 8 also includes a bound on these probabilities in the asymptotic case where $\Lambda \to \infty$. This bound, which is derived in Appendix XI, depends only on the prior probabilities, means, and variance of the nonzero-value target classes.

It is seen that the GA, GU/LA, and LA policies all perform similarly to the oracle policy in all of these metrics as SNR improves. ARAP, which does not distinguish between high- and low-value targets, performs better for low-value targets and worse for high-value targets. Note that in these plots, the GU/LA policy used only 50 local sensors in comparison to the LA policy which used 400 local sensors. The GA policy has slightly lower misclassification errors than the GU/LA and LA policies for low SNR values (SNR $<$ 10 dB).

In some cases, the ultimate goal of the mission is to use the information obtained by the sensors in

order to achieve some objective. For example, this may include dropping payloads for military purposes or for first-aid after natural disasters. Moreover, it may be the case that there are only a limited number of payloads available. While the goal of this paper is not to optimize the expected return on these payloads,the performance of the proposed policies as a function of the number of payloads is easily simulated. Define the expected return with p payloads as

$$P_T(p) = \sum_{i=1}^{p} z_{w(i)}(T), \tag{54}$$

where $z_i(T)$ is the posterior mean reward at the final stage, as defined by (48), and w is a sorting operator such that

$$z_{w(1)} \geq \cdots \geq z_{w(N)} \tag{55}$$

Thus payloads are assigned to the p locations with highest expected importance. Fig. 9 compares $P_T(p)$ as a function of p and policy with SNR = 8 dB. The GA, GU/LA, and LA policies quickly converge to the oracle policy, followed by ARAP. However, the LA policy requires 5 times more local sensors and has larger signal estimation error than the GU/LA policy.

VI. CONCLUSION

The presented policies generalized previous formulations for adaptive target search [3,4,8] to include value of information when targets have varying mission importance. New policies are proposed that are able to simultaneously locate, classify and estimate a sparse number of targets embedded in a wide area. These policies approximately optimize an objective function that includes mission importance. We also considered multiple sensor models, which include global uniform sensors, global adaptive sensors, and local adaptive sensors.

Theoretical upper bounds on the performance of adaptive policies indicate that there are many regimes in which incorporating mission importance leads to significant performance gains. While the oracle policies are not achievable in practice, the proposed GA policy performs nearly as well as the oracle when SNR is sufficiently high. In the scenario where a globally adaptive sensor is not available, the LA and GU/LA policies also nearly achieve oracle performance given enough local sensors. However, the GU/LA policy generally performs better than the LA policy when there are only a few local sensors and it requires an order of magnitude fewer local sensors to achieve within 3 dB of the maximum performance. Finally, all of the proposed policies are shown to perform well in not only minimizing the objective function, but also in reducing mean-squared estimation error, reducing misclassification probability, and increasing expected return from payloads.

Future research directions will (a) compare and contrast this formulation to one which directly optimizes the expected return rather than the proposed objective function; (b) include time-varying mission-value that can model stale information and dynamic targets; and (c) develop additional analytical results, such as performance bounds that depend directly on the proposed policies rather than the oracle policies, as well as convergence rates as a function of the scene size and number of sensing stages.

APPENDIX I
PROPOSITION 2

Proof: Given the observation history $\boldsymbol{Y}(t-1)$, the posterior distribution $p_i^{(c)}(t)$ will only depend on $y_i(t)$. $\hat{C}_i(t)$ (defined in (16)) is a function of $p_i^{(c)}(t)$, therefore also only depends on $y_i(t)$. For brevity, $\boldsymbol{Y}(t-1)$ is left out in the following deductions, but it should be noted that all quantities depend on $\boldsymbol{Y}(t-1)$.

Define $\mathrm{y}_i^{(c)}(t)$ as the decision region that will lead to $\hat{C}_i(t) = c$:

$$\mathrm{y}_i^{(c)}(t) = \{y_i(t) \in \mathbb{R} | \hat{C}_i(t) = c\} \tag{56}$$

The probability distribution of $y_i(t)$ is given by (12) and we denote its density as $f(y_i(t) \mid C_i)$. Given Assumption 1, the variance of $y_i(t)$ is equal for all $c > 1$. Denote the variance for $c = 1$ and $c > 1$ respectively as:

$$\begin{aligned} \Sigma_i^0(t) =& \nu^2/\lambda_i(t-1) \\ \Sigma_i(t) =& \sigma_i^2(t-1) + \nu^2/\lambda_i(t-1) \end{aligned} \tag{57}$$

Denote $m_c = \hat{X}_i^{(c)}(t-1)$, and rank the elements with sequence $\pi(c)$ such that:

$$0 = m_1 = m_{\pi(1)} < m_{\pi(2)} < m_{\pi(3)} < \cdots \tag{58}$$

As a consequence of the modification to the MAP classifier in (16), we have the following simplification.

Lemma 2 The decision regions $\mathrm{y}_i^{(c)}(t)$ are simple intervals ordered in the same way as in (58):

$$(-\infty, a_i^{\pi(1)}(t)],\ (a_i^{\pi(1)}(t), a_i^{\pi(2)}(t)], \cdots \tag{59}$$

and the posterior probability at a boundary $a_i^{\pi(c)}(t)$ is the same for the two adjacent classes:

$$\begin{aligned} & p_i^{(\pi(c))}(t-1) f(a_i^{\pi(c)} | C_i = \pi(c)) \\ =& p_i^{(\pi(c+1))}(t-1) f(a_i^{\pi(c)} | C_i = \pi(c+1)) \end{aligned} \tag{60}$$

Proof: First compare the posterior probability when $c \neq 1$:

$$\arg \max_{d \in \mathcal{C}, d \neq 1} p_i^{(d)}(t)$$

plug in (11) and ignore the constant in denominitor

$$= \arg \max_{d \in \mathcal{C}, d \neq 1} p_i^{(d)}(t-1) f(y_i(t)|C_i = c)$$

plug in (12) and formula for normal distribution

$$= \arg \max_{d \in \mathcal{C}, d \neq 1} \frac{p_i^{(d)}(t-1)}{\sqrt{\Sigma_i(t)}} \exp\left\{ -\frac{(y_i(t) - m_d)^2}{2\Sigma_i(t)} \right\}$$

take logrithm

$$= \arg \max_{d \in \mathcal{C}, d \neq 1} \left(\log \frac{p_i^{(d)}(t-1)}{\sqrt{\Sigma_i(t)}} - \frac{(y_i(t) - m_d)^2}{2\Sigma_i(t)} \right)$$

change max to min, and ignore constant $\log \sqrt{\Sigma_i(t)}$

$$= \arg \min_{d \in \mathcal{C}, d \neq 1} \left(\frac{(y_i(t) - m_d)^2}{2\Sigma_i(t)} - \log p_i^{(d)}(t-1) \right) \tag{61}$$

The functions within the minimization are quadratic in $y_i(t)$, and have the same shapes across classes d. Therefore, the minimum, i.e. $\mathrm{y}_i^{(c)}(t)$, are simple intervals, and equation (60) holds.

Now consider the case $\hat{C}_i(t) = 1$.

$$\begin{aligned}
& \hat{C}_i(t) = 1 \\
\leftrightarrow \;& 1 = \arg \max_{d \in \mathcal{C}} p_i^{(d)}(t) \\
\leftrightarrow \;& p_i^{(1)}(t) \geq p_i^{(d)}(t) \quad \forall d \in \mathcal{C} \\
\leftrightarrow \;& p_i^{(1)}(t-1) f(y_i(t)|C_i = 1) \geq p_i^{(d)}(t-1) f(y_i(t)|C_i = d)
\end{aligned}$$

plug in (12)

$$\begin{aligned}
\leftrightarrow \;& p_i^{(1)}(t-1) \frac{\exp\left\{ -\frac{y_i(t)^2}{2\Sigma_i^0(t)} \right\}}{\sqrt{2\pi\Sigma_i^0(t)}} \geq p_i^{(d)}(t-1) \frac{\exp\left\{ -\frac{(y_i(t) - m_d)^2}{2\Sigma_i(t)} \right\}}{\sqrt{2\pi\Sigma_i(t)}} \\
\leftrightarrow \;& \frac{p_i^{(1)}(t-1)}{p_i^{(d)}(t-1))} \geq \sqrt{\frac{\Sigma_i^0(t)}{\Sigma_i(t)}} \exp\left\{ \frac{y_i(t)^2}{2\Sigma_i^0(t)} - \frac{(y_i(t) - m_d)^2}{2\Sigma_i(t)} \right\}
\end{aligned}$$

take logrithm

$$\begin{aligned}\leftrightarrow\ & \frac{\sigma_i^2(t)}{2\Sigma_i^0(t)\Sigma_i(t)}y_i(t)^2+\frac{m_d}{\Sigma_i(t)}y_i(t)-\frac{\mu_d^2}{2\Sigma_i(t)}\\ &\le\frac{1}{2}\log\frac{\Sigma_i(t)}{\Sigma_i^0(t)}+\log\frac{p_i^{(1)}(t-1)}{p_i^{(d)}(t-1)}\quad\forall d\in\mathcal{C}\end{aligned}$$

Solve above inequality and recall the definition in (16):

$$\mathrm{y}_i^{(1)}(t)=(-\infty,a_i^{\pi(1)}(t)] \tag{62}$$

where

$$a_i^{\pi(1)}(t)=\begin{cases}0 & \text{if } A\le 0\\ -\frac{\Sigma_i^0(t)m_{\pi(2)}}{\sigma_i^2(t)}+\sqrt{A} & \text{if } A\ge 0\end{cases}$$

$$A=\frac{\Sigma_i(t)\Sigma_i^0(t)}{\sigma_i^2(t)}\left(\frac{m_{\pi(2)}^2}{\sigma_i^2(t)}-2\log\frac{p_i^{(\pi(2))}(t-1)}{p_i^{(\pi(1))}(t-1)}+\log\frac{\Sigma_i(t)}{\Sigma_i^0(t)}\right)$$

When $A<0$, the problem becomes misclassification problem over class 1 targets is 0.5, and the problem becomes trivial. Here we consider the case $A\ge 0$, in which case Equation (60) also holds.

To sum up two cases, Equation (60) holds. ■

By definition (17),

$$\begin{aligned}q_i(t)=&\Pr\left(\hat{C}_i(t)\neq C_i|\boldsymbol{Y}(t-1)\right)\\ =&1-\Pr\left(\hat{C}_i(t)=C_i|\boldsymbol{Y}(t-1)\right)\\ =&1-\sum_{c\in\mathcal{C}}p_i^{(c)}(t-1)\int_{\mathrm{y}_i^{(c)}(t)}f(y_i(t)|C_i=c)dy_i(t)\\ =&1-\tilde{q}_i(t)\end{aligned} \tag{63}$$

Where

$$\tilde{q}_i(t)=\sum_{c\in\mathcal{C}}p_i^{(c)}(t-1)\int_{\mathrm{y}_i^{(c)}(t)}f(y_i(t)|C_i=c)dy_i(t) \tag{64}$$

is the probability of correct classification.

Denote $\Phi(\cdot)$ as standard normal cumulative probability function, $\Phi'(\cdot)$ as the standard normal probability density function. Further define normalized boundaries of $\mathrm{y}_i^{\pi(c)}(t)$:

$$b_r^{\pi(1)}=\frac{a_i^{\pi(1)}(t)}{\sqrt{\Sigma_i^0(t)}},\quad\text{and} \tag{65}$$

$$b_l^{\pi(d)}=\frac{a_i^{\pi(d-1)}(t)-m_{\pi(d)}}{\sqrt{\Sigma_i(t)}},\quad b_r^{\pi(d)}=\frac{a_i^{\pi(d)}(t)-m_{\pi(d)}}{\sqrt{\Sigma_i(t)}}$$

for $d > 1$. Then

$$\begin{aligned}
&\tilde{q}_i(t) = \sum_{c\in\mathcal{C}} p_i^{(c)}(t-1) \int_{\mathrm{y}_i^{(c)}(t)} f(y_i(t)|C_i = c)dy_i(t) \\
=&p_i^{(\pi(1))}(t-1)\Phi(b_r^{\pi(1)}) \\
+&\sum_{k>1} p_i^{(\pi(k))}(t-1)\left(\Phi(b_r^{\pi(k)}) - \Phi(b_l^{\pi(k)})\right)
\end{aligned}$$

Splitting the summation and rearranging terms

$$\begin{aligned}
=&p_i^{(\pi(1))}(t-1)\Phi(b_r^{\pi(1)}) - p_i^{(\pi(2))}(t-1)\Phi(b_l^{\pi(2)}) \\
+&\sum_{k>1} p_i^{(\pi(k))}(t-1)\Phi(b_r^{\pi(k)}) - p_i^{(\pi(k+1))}(t-1)\Phi(b_l^{\pi(k+1)}) \\
=&M_i^{1,r}(t) + \sum_{k>1} M_i^k(t)
\end{aligned} \tag{66}$$

$b_l^{\pi(k)}$, $b_r^{\pi(k)}$ are functions of $a_i^{\pi(c)}(t)$, which are in turn functions of $\Sigma_i(t)$ and $\Sigma_i^0(t)$. Since $\Sigma_i(t) = \Sigma_i^0(t) + \sigma_i^2(t)$, we can conclude that $M_i^{1,r}(t)$ and $M_i^k(t)$ depend on $\lambda_i(t)$ through $\Sigma_i^0(t)$.

$$\begin{aligned}
&\frac{dM_i^{1,r}(t)}{d\Sigma_i^0(t)} \\
=&p_i^{(\pi(1))}(t-1)\Phi'(b_r^{\pi(1)})\left(-\frac{b_r^{\pi(1)}}{2\Sigma_i^0(t)} + \frac{1}{\sqrt{\Sigma_i^0(t)}}\frac{da_i^{\pi(1)}(t)}{d\Sigma_i^0(t)}\right) \\
&- p_i^{(\pi(2))}(t-1)\Phi'(b_l^{\pi(2)})\left(-\frac{b_l^{\pi(2)}}{2\Sigma_i(t)} + \frac{1}{\sqrt{\Sigma_i(t)}}\frac{da_i^{\pi(1)}(t)}{d\Sigma_i^0(t)}\right)
\end{aligned}$$

Plugging in $\Phi'(b_r^{\pi(1)})/\sqrt{\Sigma_i^0(t)} = f(a_i^{\pi(1)}(t)|\pi(1))$, $\Phi'(b_l^{\pi(2)})/\sqrt{\Sigma_i(t)} = f(a_i^{\pi(1)}(t)|\pi(2))$ and (60),

$$= - p_i^{(\pi(1))}(t-1)f(a_i^{\pi(1)}(t)|\pi(1))\frac{\sqrt{A}}{2\Sigma_i^0(t)\Sigma_i(t)} \leq 0 \tag{67}$$

Similarly, when $k > 1$,

$$\begin{aligned}
&\frac{dM_i^k(t)}{d\Sigma_i^0(t)} \\
=&p_i^{(\pi(k))}(t-1)\Phi'(b_r^{\pi(k)})\left(-\frac{b_r^{\pi(k)}}{2\Sigma_i(t)} + \frac{1}{\sqrt{\Sigma_i(t)}}\frac{da_i^{\pi(k)}(t)}{d\Sigma_i^0(t)}\right) \\
-&p_i^{(\pi(k+1))}(t-1)\Phi'(b_l^{\pi(k+1)})\left(-\frac{b_l^{\pi(k+1)}}{2\Sigma_i(t)} + \frac{1}{\sqrt{\Sigma_i(t)}}\frac{da_i^{\pi(k)}(t)}{d\Sigma_i^0(t)}\right) \\
=&p_i^{(\pi(k))}(t-1)f(a_i^{\pi(k)}(t)|\pi(k))\frac{m_{\pi(k)} - m_{\pi(k+1)}}{2\Sigma_i(t)} \leq 0
\end{aligned} \tag{68}$$

using (60). Combining (67) and (68) with (66), $\frac{d\tilde{q}_i(t)}{d\Sigma_i^0(t)} \leq 0$. Since $\frac{d\Sigma_i^0(t)}{d\lambda_i(t)} < 0$ from (57), we have

$$\frac{d\tilde{q}_i(t)}{d\lambda_i(t)} \geq 0, \qquad \frac{dq_i(t)}{d\lambda_i(t)} \leq 0$$

■

APPENDIX II
PROOF OF PROPOSITION 3

Given Assumptions 1 and 2, the number of non-zero allocations for the oracle policy is $N - N_1$, as shown in Appendix VI. Under this condition, we can plug the oracle allocation policy (19) into the cost function (18) to yield

$$\begin{aligned} J_T(\boldsymbol{\lambda}^o) &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\sum_{i=1}^{N-N_1} \frac{h(C_{\pi(i)})}{c_0 + \frac{(\Lambda + (N-N_1)c_0)\sqrt{h(C_{\pi(i)})}}{\sum_{j=1}^{N}\sqrt{h(C_{\pi(j)})}} - c_0} \right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\sum_{i=1}^{N-N_1} \sqrt{h(C_{\pi(i)})} \left(\frac{\Lambda + (N-N_1)c_0}{\sum_{j=1}^{N-N_1}\sqrt{h(C_{\pi(j)})}} \right)^{-1} \right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\left(\frac{1}{\Lambda + (N-N_1)c_0} \right) \left(\sum_{i=1}^{N-N_1} \sqrt{h(C_{\pi(i)})} \right)^2 \right], \end{aligned} \tag{69}$$

We can further simplify the expression by noting that $h(1) = 0$ (i.e. for the zero-value class) so that we can drop the permutation operator:

$$\begin{aligned} J_T(\boldsymbol{\lambda}^o) &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\left(\frac{1}{\Lambda + (N-N_1)c_0} \right) \left(\sum_{i=1}^{N} \sqrt{h(C_i)} \right)^2 \right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\left(\frac{1}{\Lambda + (N-N_1)c_0} \right) \left(\sum_{c=1}^{|\mathcal{C}|} N_c \sqrt{h(c)} \right)^2 \right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{N}} \left[\left(\frac{1}{\Lambda + (N-N_1)c_0} \right) \left(\sum_{c=2}^{|\mathcal{C}|} N_c \sqrt{h(c)} \right)^2 \right], \end{aligned} \tag{70}$$

where the last equality once again uses $h(1) = 0$ and noting that $J_T(\boldsymbol{\lambda}^o)$ is only random through the number of targets in each class $\boldsymbol{N} = \{N_c\}_{c\in\mathcal{C}} \sim \text{Multinomial}(N, \{p_c\}_{c\in\mathcal{C}})$. We note the following properties of Multinomial distributions which will be useful in further simplifying (70):

$$\left\{N_2, N_3, \ldots, N_{|\mathcal{C}|}\right\} | N_1 \sim \text{Multinomial}\left(N - N_1, \{\tilde{p}_c\}_{c=2}^{|\mathcal{C}|}\right) \tag{71}$$

$$\mathbb{E}\left[N_i N_j | N_1\right] = \begin{cases} (N-N_1)\tilde{p}_i(1-\tilde{p}_i) + (N-N_1)^2\tilde{p}_i^2, & i=j \\ -(N-N_1)\tilde{p}_i\tilde{p}_j + (N-N_1)^2\tilde{p}_i\tilde{p}_j, & i\neq j \end{cases} \tag{72}$$
$$\forall i,j \in \{2,3,\ldots,|\mathcal{C}|\}$$

where $\tilde{p}_c \stackrel{\triangle}{=} p_c/(1-p_1)$. Note that $\{\tilde{p}_c\}_{c=2}^{|\mathcal{C}|}$ defines a probability distribution on C_i conditioned on $C_i > 1$. Using iterated expectation over (70) and the properties above:

$$\begin{aligned} &J_t(\boldsymbol{\lambda}^o) \\ &= \mathbb{E}_{N_1}\left[\left(\frac{\nu^2}{\Lambda+(N-N_1)c_0}\right)\mathbb{E}_{\boldsymbol{N}\setminus N_1|N_1}\left[\left(\sum_{c=2}^{|\mathcal{C}|} N_c\sqrt{h(c)}\right)^2\right]\right] \end{aligned} \tag{73}$$

where $\boldsymbol{N}\setminus N_1 = \left\{N_2,\ldots,N_{|\mathcal{C}|}\right\}$. Expanding the inner expectation

$$\begin{aligned} &\mathbb{E}_{\boldsymbol{N}\setminus N_1|N_1}\left[\left(\sum_{c=2}^{|\mathcal{C}|} N_c\sqrt{h(c)}\right)^2\right] \\ &= \mathbb{E}_{\boldsymbol{N}\setminus N_1|N_1}\left[\sum_{c=2}^{|\mathcal{C}|} N_c^2 h(c) + 2\sum_{c=2}^{|\mathcal{C}|}\sum_{d=c+1}^{|\mathcal{C}|} N_c N_d\sqrt{h(c)h(d)}\right] \\ &= \sum_{c=2}^{|\mathcal{C}|}\left((N-N_1)\tilde{p}_c(1-\tilde{p}_c) + (N-N_1)^2\tilde{p}_c^2\right)h(c) \\ &\quad + 2\sum_{c=2}^{|\mathcal{C}|}\sum_{d=c+1}^{|\mathcal{C}|}\left(-(N-N_1)\tilde{p}_c\tilde{p}_d + (N-N_1)^2\tilde{p}_c\tilde{p}_d\right)\sqrt{h(c)h(d)} \\ &= (N-N_1)a_1 + (N-N_1)^2 a_2 \end{aligned} \tag{74}$$

where

$$a_1 = \sum_{c=2}^{|\mathcal{C}|}\tilde{p}_c(1-\tilde{p}_c)h(c) - 2\sum_{c=2}^{|\mathcal{C}|}\sum_{d=c+1}^{|\mathcal{C}|}\tilde{p}_c\tilde{p}_d\sqrt{h(c)h(d)} \tag{75}$$

$$a_2 = \sum_{c=2}^{|\mathcal{C}|}\tilde{p}_c^2 h(c) + 2\sum_{c=2}^{|\mathcal{C}|}\sum_{d=c+1}^{|\mathcal{C}|}\tilde{p}_c\tilde{p}_d\sqrt{h(c)h(d)} \tag{76}$$

With some simple algebraic manipulation, it can easily be shown that

$$a_2 = \left(\sum_{c=2}^{|\mathcal{C}|}\tilde{p}_c\sqrt{h(c)}\right)^2 = m_{h,1}^2 \tag{77}$$

and

$$a_1 = \sum_{c=2}^{|\mathcal{C}|}\tilde{p}_c h(c) - a_2 = m_{h,2} - m_{h,1}^2 \tag{78}$$

Using (74) in (73) yields an expression for the oracle cost that is simply an expectation over a single random parameter $N_1 \sim \text{Binomial}(N, p_1)$:

$$J_t(\boldsymbol{\lambda}^o) = \nu^2 \mathbb{E}_{N_1}\left[\frac{(N-N_1)a_1 + (N-N_1)^2 a_2}{\Lambda + (N-N_1)c_0}\right] \tag{79}$$

To complete the proof, we use Lemma 3 below to bound the expectation over N_1.

Lemma 3 Let $f : [0, N] \to \mathbf{R}$ be a function of the form

$$f(x) = (e_1 x + e_2 x^2)/(\Lambda + x c_0) \tag{80}$$

where (e_1, e_2) satisfy $\Lambda e_2 \geq c_0 e_1$. Then for an even integer $d \geq 2$ and a random variable $X \in [0, N]$, we have

$$\mathbb{E}\left[f(X)\right] \geq \mathbb{E}\left[P^f_{d-1}(x)\right] \tag{81}$$

$$\mathbb{E}\left[f(X)\right] \leq \mathbb{E}\left[P^f_{d-1}(x)\right] + \frac{f^{(d)}(0)\mathbb{E}\left[(X-\mu_X^{(1)})^d\right]}{d!} \tag{82}$$

$$\mathbb{E}\left[P^f_{d-1}(x)\right] = \sum_{w=0}^{d-1} \frac{f^{(w)}(\mu_X^{(1)})\mathbb{E}\left[(X-\mu_X^{(1)})^w\right]}{w!} \tag{83}$$

$$f^{(w)}(x) = \frac{(-1)^w(\Lambda e_2 - c_0 e_1)\Lambda c_0^{w-2} w!}{(\Lambda + c_0 x)^{w+1}}, \quad w \geq 2 \tag{84}$$

where $f^{(w)}(x)$ is the w-th derivative of $f(x)$ and $\mu_X^{(w)}$ is the w-th central moment.

Proof: The proof is given in Appendix III. ■

Note that we can rewrite (79) as

$$J_t(\boldsymbol{\lambda}^o) = \nu^2 \mathbb{E}\left[f(N-N_1)\right] \tag{85}$$

with $e_1 = a_1$ and $e_2 = a_2$. Using Assumption 2, we have

$$\Lambda \geq c_0 \Lambda_1 = c_0\left(\frac{m_{h,2}}{m_{h,1}^2} - 1\right) = c_0 a_1/a_2 \tag{86}$$

so that

$$a_2\Lambda - a_1 c_0 \geq a_2(c_0 a_1/a_2) - a_1 c_0 = 0 \tag{87}$$

Therefore, we can apply Lemma 3 with $X = N - N_1$ and $d = 2$:

$$\begin{aligned} J_t(\boldsymbol{\lambda}^o) &\geq \nu^2\left(f(\mu_X^{(0)}) + f^{(1)}(\mu_X^{(0)})\mathbb{E}\left[X - \mu_X^{(0)}\right]/2\right) \\ &= \nu^2 f(N(1-p_1)) \\ &= \nu^2\left(\frac{(N-Np_1)a_1 + (N-Np_1)^2 a_2}{\Lambda + (N-Np_1)c_0}\right) \end{aligned} \tag{88}$$

where we used the fact that $\mathbb{E}\left[X-\mu_X^{(0)}\right]=0$ and $\mu_X^{(0)}=N(1-p_1)$. The result follows from simple factorization of (88) and using the definitions in (78) and (77). Note that we can get a tighter bound by applying Lemma 3 with $d>2$. To derive the upper bound, we apply (82) from 3 with $d=2$, where we notice that the first term is just the lower bound and the second term is given by

$$\begin{aligned} & f^{(2)}(0)\mathbb{E}\left[(X-\mu_X^{(1)})^2\right]/d! \\ & \quad = \frac{(\Lambda m_{h,1}^2 - c_0(m_{h,2}-m_{h,1}^2))Np_1(1-p_1)}{\Lambda^2} \\ & \quad = \frac{((\Lambda+c_0)m_{h,1}^2 - c_0 m_{h,2}))Np_1(1-p_1)}{\Lambda^2} \end{aligned} \tag{89}$$

where the expectation is just the variance of the random variable $N-N_1$, which is $Np_1(1-p_1)$.

APPENDIX III
PROOF OF LEMMA 3

First note that the first two derivatives of $f(x)$ are given by

$$f^{(1)}(x) = \Lambda(e_1+2e_2x) + e_2c_0x^2/(\Lambda+c_0x)^2 \tag{90}$$

$$f^{(2)}(x) = 2\Lambda(\Lambda e_2 - e_1c_0)/(\Lambda+c_0x)^3 \tag{91}$$

Only the denominator of $f^{(2)}(x)$ depends on x, from which the relationship (84) can easily be derived. Define the d-th order Taylor series approximation to f around point c to be

$$P_d^f(x) = \sum_{w=0}^{d} \frac{f^{(w)}(c)}{w!}(x-c)^w \tag{92}$$

Using Taylor's theorem, we have

$$f(x) = P_{d-1}^f(x) + \frac{f^{(d)}(b)}{d!}(x-c)^d \tag{93}$$

for some $b\in[0,N]$ (i.e., the domain of f). In particular, when d is even, we have both $f^{(d)}(x)\geq 0$ and $(x-c)^d\geq 0$ for all $x\in[0,N]$. Therefore

$$f(x) \geq P_{d-1}^f(x) \tag{94}$$

Moreover, using the corollary of Taylor's theorem, we have

$$\left|f(x)-P_{d-1}^f(x)\right| \leq \frac{\max\limits_{b\in[0,N]}|f^{(d)}(b)|}{d!}|x-c|^d \tag{95}$$

From (84), it is clear the maximum of $|f^{(d)}(b)|$ occurs when $b = 0$. Moreover, since d is even $|x-c|^d = (x-c)^d$ and $f^{(d)}(0) \geq 0$. Since $f(x) > P^f_{d-1}(x)$, we can conclude that

$$\begin{aligned} & f(x) - P^f_{d-1}(x) \leq \frac{f^{(d)}(0)}{d!}(x-c)^d \\ \Leftrightarrow \qquad & f(x) \leq P^f_{d-1}(x) + \frac{f^{(d)}(0)}{d!}(x-c)^d \end{aligned} \tag{96}$$

(81) and (82) result from applying linearity of expectation to 94 and 96 with Taylor series expansion around $c = \mathbb{E}\left[X - \mathbb{E}\left[X\right]\right]$.

APPENDIX IV
PROOF OF PROPOSITION 1

We start by using linearity of the summation operators to rewrite the cost (14) as

$$J_T(\boldsymbol{\lambda}) = \mathbb{E}\left[\sum_{c\in\mathcal{C}} h(c)\left(\sum_{i=1}^{N} I_i^{(c)}(X_i - \hat{X}_i^{(c)}(T))^2\right)\right], \tag{97}$$

To simplify (14), it is useful to note that each term of the inner summation is non-zero only when $C_i = c$. Moreover, using iterated expectation and the linearity of expectation, we have $J_T(\boldsymbol{\lambda}) =$

$$\sum_{c\in\mathcal{C}} h(c)\mathbb{E}_{\boldsymbol{Y}(T)}\left[\mathbb{E}\left[\sum_{i=1}^{N} I_i^{(c)}(X_i - \hat{X}_i^{(c)}(T))^2 \Big| C_i = c, \boldsymbol{Y}(T)\right]\right]. \tag{98}$$

Given $C_i = c$, $I_i^{(c)}$ is deterministic and hence independent of the other terms. Using this property, (6), and the definition of variance, we have

$$\begin{aligned} J_T(\boldsymbol{\lambda}) &= \sum_{c\in\mathcal{C}} h(c)\mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N} p_i^{(c)}(T)(\sigma_i^{(c)}(T))^2\right] \\ &= \mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N}\sum_{c\in\mathcal{C}} h(c)p_i^{(c)}(T)(\sigma_i^{(c)}(T))^2\right], \end{aligned} \tag{99}$$

where $p_i^{(c)}(T) = \Pr(C_i = c|\boldsymbol{Y}(T))$ and $(\sigma_i^{(c)}(T))^2 = \mathrm{var}\left[X_i|C_i = c, \boldsymbol{Y}(T)\right]$.

When conditioned on $C_i = c$, we know that X_i is Gaussian by definition. It can be shown that when X_i is Gaussian and $\boldsymbol{\lambda}(t-1)$ is a deterministic function of previous measurements, then any new measurements $\boldsymbol{y}(t)$ are also Gaussian. This leads to the following simple recursion on the conditional variances:

$$\frac{1}{(\sigma_i^{(c)}(T+1))^2} = \frac{1}{(\sigma_i^{(c)}(T))^2} + \frac{\lambda_i(T)}{\nu^2} = \frac{1}{\sigma_c^2} + \frac{\sum_{t=0}^{T}\lambda_i(t)}{\nu^2}, \tag{100}$$

so that

$$(\sigma_i^{(c)}(T+1))^2 = \nu^2\left[\frac{\nu^2}{\sigma_c^2} + \sum_{t=0}^{T}\lambda_i(t)\right]^{-1} = \nu^2\left[\frac{\nu^2}{\sigma_c^2} + \overline{\lambda_i}\right]^{-1} \tag{101}$$

where $\overline{\lambda_i} = \sum_{t=0}^{T-1} \lambda_i(t)$. Plugging this into (99), we get

$$\begin{aligned} J_T(\boldsymbol{\lambda}) &= \mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N}\sum_{c\in\mathcal{C}} h(c)p_i^{(c)}(T)\sigma_i^2(T)\right] \\ &= \nu^2\mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N}\sum_{c\in\mathcal{C}} \frac{h(c)p_i^{(c)}(T)}{\nu^2/\sigma_c^2+\overline{\lambda_i}}\right] \end{aligned} \tag{102}$$

When Assumption 1 holds, we can simplify this further, so that

$$\begin{aligned} J_T(\boldsymbol{\lambda}) &= \nu^2\mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N} \frac{\sum_{c\in\mathcal{C}} h(c)p_i^{(c)}(T)}{\nu^2/\sigma_0^2+\overline{\lambda_i}}\right] \\ &= \nu^2\mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N} \frac{z_i(T)}{\nu^2/\sigma_0^2+\overline{\lambda_i}}\right]. \end{aligned} \tag{103}$$

where $z_i(T) = \sum_{c\in\mathcal{C}} h(c)p_i^{(c)}(T)$. Using iterated expecation and Lemma 4 (below), we get the final result

$$\begin{aligned} J_T(\boldsymbol{\lambda}) &= \nu^2\mathbb{E}_{\boldsymbol{Y}(T-1)}\left[\mathbb{E}_{\boldsymbol{y}(T)|\boldsymbol{Y}(T-1)}\left[\sum_{i=1}^{N} \frac{z_i(T)}{\nu^2/\sigma_0^2+\overline{\lambda_i}}\right]\right], \\ &= \nu^2\mathbb{E}_{\boldsymbol{Y}(T-1)}\left[\sum_{i=1}^{N} \frac{\mathbb{E}_{\boldsymbol{y}(T)|\boldsymbol{Y}(T-1)}\left[z_i(T)\right]}{\nu^2/\sigma_0^2+\overline{\lambda_i}}\right], \\ &= \nu^2\mathbb{E}_{\boldsymbol{Y}(T-1)}\left[\sum_{i=1}^{N} \frac{z_i(T-1)}{\nu^2/\sigma_0^2+\overline{\lambda_i}}\right]. \end{aligned} \tag{104}$$

Note that using (101), we can rewrite an equivalent result in terms of the final-stage allocations $\boldsymbol{\lambda}(T-1)$ as

$$J_T(\boldsymbol{\lambda}) = \nu^2\mathbb{E}_{\boldsymbol{Y}(T-1)}\left[\sum_{i=1}^{N} \frac{z_i(T-1)}{\nu^2/\sigma_i^2(T-1)+\lambda_i(T-1)}\right]. \tag{105}$$

The derivation of the final expressions in (15) without Assumption 1 is entirely analogous.

Lemma 4

$$\mathbb{E}_{\boldsymbol{y}(T)|\boldsymbol{Y}(T-1)}\left[z_i(T)\right] = z_i(T-1) \tag{106}$$

Proof: Note that

$$\begin{aligned} \Pr(\boldsymbol{C}|\boldsymbol{Y}(T)) &= \Pr(\boldsymbol{C}|\boldsymbol{Y}(T-1),\boldsymbol{y}(T)) \\ &= \frac{\Pr\left(\boldsymbol{y}(t)|\boldsymbol{C},\boldsymbol{Y}(t-1)\right)\Pr\left(\boldsymbol{C}|\boldsymbol{Y}(T-1)\right)}{\Pr\left(\boldsymbol{y}(t)|\boldsymbol{Y}(T-1)\right)} \end{aligned} \tag{107}$$

Then

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{y}(T)|\boldsymbol{Y}(T-1)}\left[\Pr(\boldsymbol{C}|\boldsymbol{Y}(T))\right] \\
&= \int \frac{\Pr\left(\boldsymbol{y}(t)|\boldsymbol{C},\boldsymbol{Y}(t-1)\right)\Pr\left(\boldsymbol{C}|\boldsymbol{Y}(T-1)\right)}{\Pr\left(\boldsymbol{y}(t)|\boldsymbol{Y}(T-1)\right)} \\
&\qquad \cdot \Pr\left(\boldsymbol{y}(t)|\boldsymbol{Y}(T-1)\right) d\boldsymbol{y}(t) \\
&= \int \Pr\left(\boldsymbol{C}|\boldsymbol{Y}(T-1)\right)\Pr\left(\boldsymbol{y}(t)|\boldsymbol{C},\boldsymbol{Y}(t-1)\right) d\boldsymbol{y}(t) \\
&= \Pr\left(\boldsymbol{C}|\boldsymbol{Y}(T-1)\right) \int \Pr\left(\boldsymbol{y}(t)|\boldsymbol{C},\boldsymbol{Y}(t-1)\right) d\boldsymbol{y}(t) \\
&= \Pr\left(\boldsymbol{C}|\boldsymbol{Y}(T-1)\right)
\end{aligned} \tag{108}$$

Note that $f\left(\boldsymbol{C}|\boldsymbol{Y}(T)\right)$ is the multivariate extension of $p_i^{(c)}(T) = \Pr(C_i = c|\boldsymbol{Y}(T))$, from which we can conclude

$$\mathbb{E}_{\boldsymbol{y}(T)|\boldsymbol{Y}(T-1)}\left[p_i^{(c)}(T)\right] = p_i^{(c)}(T-1) \tag{109}$$

Using (109) in the definition of $z_i(T)$ yields the result. ■

APPENDIX V
PROOF OF LEMMA 1

Applying the law of total expectation to (15)

$$\begin{aligned}
J_T(\boldsymbol{\lambda}) &= \nu^2 \mathbb{E}_{\boldsymbol{Y}(T)}\left[\sum_{i=1}^{N}\sum_{c=2}^{|\mathcal{C}|}\frac{h(c)p_i^{(c)}(T-1)}{\nu^2/\sigma_c^2 + \overline{\lambda_i}(T-1)}\right] \\
&= \nu^2 \mathbb{E}_{\boldsymbol{C}}\left[\mathbb{E}_{\boldsymbol{Y}(T-1)|\boldsymbol{C}}\left[\sum_{i=1}^{N}\sum_{c=2}^{|\mathcal{C}|}\frac{h(c)p_i^{(c)}(T-1)}{\nu^2/\sigma_c^2 + \overline{\lambda_i}(T-1)}\right]\right]
\end{aligned} \tag{110}$$

Since $\lambda_i(t)$ is a function only of $\boldsymbol{C}$, the denominator conditioned on $\boldsymbol{C}$ is deterministic so that

$$J_T(\boldsymbol{\lambda}) = \nu^2 \mathbb{E}_{\boldsymbol{C}}\left[\sum_{i=1}^{N}\sum_{c=2}^{|\mathcal{C}|}\frac{\mathbb{E}_{\boldsymbol{Y}(T-1)|\boldsymbol{C}}\left[h(c)p_i^{(c)}(T-1)\right]}{\nu^2/\sigma_c^2 + \overline{\lambda_i}(T-1)}\right] \tag{111}$$

Moreover, we have

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{Y}(t)|\boldsymbol{C}}\left[\sum_{c\in\mathcal{C}} h(c)p_i^{(c)}(t)\right] &= \mathbb{E}_{\boldsymbol{Y}(t)|\boldsymbol{C}}\left[\sum_{c\in\mathcal{C}} h(c)\mathbb{E}_{C_i}\left[I_i^{(c)}(t)\right]\right] \\
&= \mathbb{E}_{\boldsymbol{Y}(t)|\boldsymbol{C}}\left[\sum_{c\in\mathcal{C}} h(c)I_i^{(c)}(t)\right] = h(C_i)
\end{aligned} \tag{112}$$

Plugging (112) into (110) yields

$$J_T(\boldsymbol{\lambda}) = \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\sum_{i=1}^{N} \sum_{c=2}^{|\mathcal{C}|} \frac{h(C_i)}{\nu^2/\sigma_c^2 + \overline{\lambda}_i (T-1)} \right]. \tag{113}$$

Note that Assumption 1 further simplifies the denominator.

APPENDIX VI
DERIVATION OF (19)

The solution here follows [8] and begins by letting $c_i = \nu^2/\sigma_{C_i}^2$. Then, define $g(k)$ to be the monotonically non-decreasing function of $k = 0, \ldots, N$ with $g(0) = 0$,

$$g(k) = \frac{c_{\pi(k+1)}}{\sqrt{h(C_{\pi(k+1)})}} \sum_{i=1}^{k} \sqrt{h(C_{\pi(i)})} - \sum_{i=1}^{k} c_{\pi(i)}, \tag{114}$$

for $k = 1, \ldots, N-1$, and $g(N) = \infty$. Define k^* by the interval $(g(k-1), g(k)]$ to which the budget parameter Λ belongs. Since $g(k)$ is monotonic, the mapping from $\Lambda(t)$ to k^* is one-to-one. The optimal full oracle policy is then

$$\overline{\lambda}^o_{\pi(i)} = \left(\Lambda(t) + \sum_{j=1}^{k^*} c_{\pi(j)} \right) \frac{\sqrt{h(C_{\pi(i)})}}{\sum_{j=1}^{k^*} \sqrt{h(C_{\pi(j)})}} - c_{\pi(i)}, \tag{115}$$

when $i \leq k^*$, and zero else.

We can simplify this expression when Assumption 1 holds. In particular, from (20), we can easily see that

$$h(C_{\pi(i)}) = \begin{cases} h(|\mathcal{C}|), & i = 1, 2, \ldots, N_{|\mathcal{C}|} \\ h(|\mathcal{C}|-1), & i = N_{|\mathcal{C}|}+1, \ldots, N_{|\mathcal{C}|} + N_{|\mathcal{C}|-1} \\ \vdots, & \vdots \\ h(2), & i = N - N_1 - N_2 + 1, \ldots, N - N_1 \\ h(1) = 0, & i = N - N_1 + 1, \ldots, N \end{cases} \tag{116}$$

Inspecting (116) further, we note an important property in the case where $C_{\pi(i)} \neq C_{\pi(i+1)}$. Specifically, we have

$$\pi(i) = \sum_{c=d}^{|\mathcal{C}|} N_c \tag{117}$$

for some $d \in \{2, 3, \ldots, |\mathcal{C}|\}$. Moreover, if d is chosen to satisfy (117) and $N_c > 0$ for all $c \in \mathcal{C}$, then we have

$$\begin{aligned} h(C_{\pi(i)}) &= h(d) \\ h(C_{\pi(i)+1}) &= h(d-1) \end{aligned} \tag{118}$$

Thus, we have

$$g(k) = c_0 \left(\sum_{i=1}^{k} \frac{\sqrt{h(C_{\pi(i)})}}{\sqrt{h(C_{\pi(k+1)})}} - k \right) \tag{119}$$

where $c_0 = \nu^2/\sigma_0^2$. There are two regimes of interest to further simplify this function: (a) when $C_{\pi(k+1)} = C_{\pi(k)}$ and (b) when $C_{\pi(k+1)} \neq C_{\pi(k)}$. In the former case we note that

$$\begin{aligned} g(k) &= c_0 \left(\sum_{i=1}^{k} \frac{\sqrt{h(C_{\pi(i)})}}{\sqrt{h(C_{\pi(k)})}} - k \right) \\ &= c_0 \left(\sum_{i=1}^{k-1} \frac{\sqrt{h(C_{\pi(i)})}}{\sqrt{h(C_{\pi(k)})}} - (k-1) \right) = g(k-1) \end{aligned} \tag{120}$$

Therefore, within intervals where $C_{\pi(k)} = C_{\pi(k+1)}$, we have equal values of $g(k)$. This implies that when there is sufficient budget to search for a single target of a given class, then we should search for all targets of that class. When $C_{\pi(k+1)} \neq C_{\pi(k)}$, we note from (117) and (118) that

$$\begin{aligned} g(\pi(k)) &= c_0 \left(\frac{1}{\sqrt{h(\tilde{d}(k)-1)}} \sum_{i=1}^{k} \sqrt{h(C_{\pi(i)})} - \sum_{c=\tilde{d}(k)}^{|\mathcal{C}|} N_c \right) \\ &= c_0 \left(\frac{1}{\sqrt{h(\tilde{d}(k)-1)}} \sum_{c=\tilde{d}(k)}^{|\mathcal{C}|} N_c \sqrt{h(c)} - \sum_{c=\tilde{d}(k)}^{|\mathcal{C}|} N_c \right) \end{aligned} \tag{121}$$

where

$$\tilde{d}(k) = \left\{ d \in \{2, 3, \ldots, |\mathcal{C}|\} : \sum_{c=d}^{|\mathcal{C}|} N_c \leq k < \sum_{c=d-1}^{|\mathcal{C}|} N_c \right\} \tag{122}$$

Using (120), (121), (122), and noticing that $g(1) = 0$, we have

$$g(k) = \begin{cases} 0, & k < N_{|\mathcal{C}|}, \\ \infty, & k \geq N - N_1 \\ c_0 \sum\limits_{c=\tilde{d}(k)}^{|\mathcal{C}|} N_c \left(\dfrac{\sqrt{h(c)}}{\sqrt{h(\tilde{d}(k)-1)}} - 1 \right), & \text{else.} \end{cases} \tag{123}$$

Note that $\tilde{d}(k)$ is constant for all targets of the same class, and thus so is $g(k)$. The number of nonzero allocations k^* is given by the interval where $\Lambda \in (g(k^*-1), g(k^*)]$. We can conclude that when the oracle policy allocates any resources to a target with class c, then it must allocate to all targets with class

c. The oracle policy is then given by

$$\overline{\lambda^o_{\pi(i)}} = \begin{cases} \dfrac{(\Lambda + k^* c_0)\sqrt{h(C_{\pi(i)})}}{\sum_{j=1}^{k^*}\sqrt{h(C_{\pi(j)})}} - c_0, & i = 1, \ldots, k^* \\ 0, & \text{else} \end{cases} \tag{124}$$

We can simplify this further using Assumption 2 where we have

$$\begin{aligned} \Lambda \geq c_0\Lambda_0 &= c_0 \sum_{c=2}^{|\mathcal{C}|} N_c \sqrt{\frac{h(c)}{h(2)}} - (N - N_1) \\ &= g(N - N_1 - 1) \end{aligned} \tag{125}$$

Since $g(N - N_1) = \infty$, we have $\Lambda \in (g(N - N_1 - 1), g(N - N_1)]$ so that $k^* = N - N_1$. Plugging into (19), we have

$$\overline{\lambda^o_{\pi(i)}} = (\Lambda + (N - N_1)c_0)\left(\frac{\sqrt{h(C_i)}}{\sum_{c=2}^{|\mathcal{C}|} N_c\sqrt{h(c)}}\right) - c_0 \tag{126}$$

for $C_i > 1$ and zero else.

APPENDIX VII
PROBABILITY OF SATISFYING (24)

Assumption 2 requires that $\Lambda \geq \Lambda_{\min} \geq c_0\Lambda_0$, where Λ_0 depends on random variables $\{N_c\}_{c=1}^{|\mathcal{C}|}$. Here we show that Λ_0 is bounded from above by Λ_0' in (28) with probability converging exponentially to 1 as N increases. Hence requiring $\Lambda_{\min} \geq c_0\Lambda_0'$ implies that $\Lambda_{\min} \geq c_0\Lambda_0$ with the same probability or greater.

First we recognize Λ_0 as the sum of N i.i.d. random variables $G_1, \ldots, G_N$, where

$$G_i = \begin{cases} 0, & C_i = 1, \\ \left(\sqrt{\frac{h(C_i)}{h(2)}} - 1\right), & C_i > 1. \end{cases} \tag{127}$$

Given the ordering of classes, G_i is bounded above and below by $G_{\max} = \sqrt{h(|\mathcal{C}|)/h(2)} - 1$ and 0 respectively. The mean of G_i is

$$\mu_G = \sum_{c=2}^{|\mathcal{C}|} p_c \left(\sqrt{\frac{h(c)}{h(2)}} - 1\right) = \bar{p}_1 \left(\frac{m_{h,1}}{\sqrt{h(2)}} - 1\right). \tag{128}$$

We may now apply the Chernoff-Hoeffding bound for independent bounded random variables [17] to the probability that Λ_0 exceeds its mean $N\mu_G$ by a multiplicative factor $\alpha > 1$. Using (28), this yields the exponentially decaying bound

$$\Pr(\Lambda_0 > \alpha N\mu_G) = \Pr(\Lambda_0 > \Lambda_0') \leq \exp\left(-ND(\alpha\hat{p}_1 \| \hat{p}_1)\right), \tag{129}$$

where

$$D(p\|q) = p\log\left(\frac{p}{q}\right) + (1-p)\log\left(\frac{1-p}{1-q}\right) \tag{130}$$

is the Kullback-Leibler divergence between Bernoulli random variables, and

$$\hat{p}_1 = \frac{\mu_G}{G_{\max}} = \bar{p}_1 \frac{m_{h,1} - \sqrt{h(2)}}{\sqrt{h(|\mathcal{C}|)} - \sqrt{h(2)}}. \tag{131}$$

APPENDIX VIII
PROOF OF PROPOSITION 4

Plugging the location-only oracle allocation policy (31) into the cost function (18), we get

$$\begin{aligned} J_T(\boldsymbol{\lambda}^{lo}) &= \nu^2 \mathbb{E}_{\boldsymbol{C}}\left[\sum_{i=1}^{N} \frac{h(C_i)}{c_0 + \Lambda/(N-N_1)}\right] \\ &= \nu^2 \mathbb{E}_{\boldsymbol{N}}\left[\frac{\sum_{c=2}^{|\mathcal{C}|} N_c h(c)}{c_0 + \Lambda/(N-N_1)}\right] \end{aligned} \tag{132}$$

Once again using properties of the multinomial distribution conditioned on N_1, we have

$$\begin{aligned} J_T(\boldsymbol{\lambda}^{lo}) &= \nu^2 \mathbb{E}_{N_1}\left[\frac{\sum_{c=2}^{|\mathcal{C}|}(N-N_1)\tilde{p}_c h(c)}{c_0 + \Lambda/(N-N_1)}\right] \\ &= \left(\frac{\nu^2 \sum_{c=2}^{|\mathcal{C}|} p_c h(c)}{1-p_1}\right) \mathbb{E}_{N_1}\left[\left(\frac{(N-N_1)^2}{c_0(N-N_1)+\Lambda}\right)\right] \\ &= \nu^2 m_{h,2} \mathbb{E}_{N_1}\left[\left(\frac{(N-N_1)^2}{c_0(N-N_1)+\Lambda}\right)\right] \end{aligned} \tag{133}$$

The expression within the expecation can be evaluated using Lemma 3 with $e_2 = 1$, $e_1 = 0$, $X = N - N_1$, $\mu_X^{(1)} = N(1-p_1)$ and $d = 2$ so that

$$\begin{aligned} J_T(\boldsymbol{\lambda}^{lo}) &\geq \nu^2 m_{h,2} f(N(1-p_1)) \\ &= \nu^2 m_{h,2}\left(\frac{(N-Np_1)^2}{c_0(N-Np_1)+\Lambda}\right) \\ &= \nu^2 m_{h,2}\left(\frac{N^2(1-p_1)^2}{c_0 N(1-p_1)+\Lambda}\right) \\ &= \frac{\nu^2 N^2 (1-p_1)^2 m_{h,2}}{\Lambda + N(1-p_1)c_0}. \end{aligned} \tag{134}$$

To derive the upper bound, we apply (82) from Lemma 3 to the expecation in (132) with $d = 2$, where we notice that the first term of (82) is just the lower bound stated above and the second term is given by

$$\begin{aligned} & f^{(2)}(0)\mathbb{E}\left[(X-\mu_X^{(1)})^2\right]/d! \\ &= \frac{\Lambda^2 N p_1(1-p_1)}{\Lambda^3} \\ &= \frac{N p_1(1-p_1)}{\Lambda} \end{aligned} \tag{135}$$

where the expectation is once again the variance of the random variable $N - N_1$, which is $Np_1(1-p_1)$.

APPENDIX IX
PROOF OF PROPOSITION 5

Using Lemma 1 and (35)

$$\begin{aligned} J_T(\boldsymbol{\lambda}^u) &= \nu^2 \mathbb{E}\left[\sum_{i=1}^{N} \frac{h(C_i)}{\nu^2/\sigma_0^2 + \Lambda/N}\right] \\ &= \left(\frac{\nu^2}{\nu^2/\sigma_0^2 + \Lambda/N}\right) \mathbb{E}\left[\sum_{i=1}^{N} h(C_i)\right] \\ &= \left(\frac{\nu^2}{\nu^2/\sigma_0^2 + \Lambda/N}\right) \mathbb{E}\left[\sum_{c=2}^{|\mathcal{C}|} N_c h(c)\right] \\ &= \left(\frac{\nu^2 N}{\nu^2/\sigma_0^2 + \Lambda/N}\right) \sum_{c=2}^{|\mathcal{C}|} p_c h(c) = \frac{\nu^2 N \bar{p}_1 m_{h,2}}{\nu^2/\sigma_0^2 + \Lambda/N} \end{aligned} \tag{136}$$

where the second-to-last equality can be evaluated noting that $\boldsymbol{N} \sim \text{Multinomial}(N, \{p_c\}_{c\in\mathcal{C}})$.

APPENDIX X
FULL-ORACLE COST WITH GENERAL PRIOR TARGET VARIANCES

Plugging (19) into the cost function (18) yields

$$\begin{aligned} J_T(\boldsymbol{\lambda}^o) &= \nu^2 \mathbb{E}_{\boldsymbol{C}}\left[\sum_{i=1}^{N} \frac{h(C_{\pi(i)})}{\dfrac{\nu^2}{\sigma^2_{C_{\pi(i)}}} + \dfrac{\left(\Lambda + \nu^2 \sum\limits_{j=1}^{k^*} \sigma^{-2}_{C_{\pi(j)}}\right)\sqrt{h(C_{\pi(i)})}}{\sum_{j=1}^{k^*}\sqrt{h(C_{\pi(j)})}} - \dfrac{\nu^2}{\sigma^2_{C_{\pi(i)}}}}\right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{C}}\left[\left(\Lambda + \nu^2 \sum_{j=1}^{k^*} \sigma^{-2}_{C_{\pi(j)}}\right)^{-1} \left(\sum_{i=1}^{N} \sqrt{h(C_{\pi(i)})}\right)^2\right], \end{aligned} \tag{137}$$

As in the equal-variance case, given sufficent SNR (similar to Assumption 2), the number of non-zero allocations for the oracle policy is $N - N_1$ and we can further simplify the expression by noting that $h(1) = 0$ (i.e. for the zero-value class) so that we can drop the permutation operator and replace the summations over the classes and number of targets in each class:

$$\begin{aligned} J_T(\boldsymbol{\lambda}^o) &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\left(\Lambda + \nu^2 \sum_{c=2}^{|\mathcal{C}|} N_c \sigma_c^{-2} \right)^{-1} \left(\sum_{i=1}^{N} \sqrt{h(C_i)} \right)^2 \right], \\ &= \nu^2 \mathbb{E}_{\boldsymbol{N}} \left[\left(\Lambda + \nu^2 \sum_{c=2}^{|\mathcal{C}|} N_c \sigma_c^{-2} \right)^{-1} \left(\sum_{c=1}^{|\mathcal{C}|} N_c \sqrt{h(c)} \right)^2 \right], \end{aligned} \tag{138}$$

where the last equality once again uses $h(1) = 0$ and noting that $J_T(\boldsymbol{\lambda}^o)$ is only random through the number of targets in each class $\boldsymbol{N} = \{N_c\}_{c \in \mathcal{C}} \sim \text{Multinomial}(N, \{p_c\}_{c \in \mathcal{C}})$.

As opposed to the previous specialized case of equal prior target variance, there is no simple way to analytically evaluate this expectation. Nevertheless, it is simple to evaluate the expectation through Monte Carlo approximation by taking samples from $\boldsymbol{N}$. In this way, we can evaluate the bounds to compare to the equal-variance case. Note that only the first quantity in (138) is different in comparison to the equal-variance case, and it only depends on the parameters through the prior target variances and the number of targets in each class.

To complete the analysis, we also include the costs of the location-only and uniform policies in the general case. In particular, following Appendices VIII and IX with appropriate modifications, we have

$$\begin{aligned} J_T(\boldsymbol{\lambda}^{lo}) &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\sum_{i=1}^{N} \frac{h(C_i)}{\nu^2/\sigma_{C_i}^2 + \Lambda/(N - N_1)} \right] \\ &= \nu^2 \mathbb{E}_{\boldsymbol{N}} \left[\sum_{c=2}^{|\mathcal{C}|} \frac{N_c h(c)}{\nu^2/\sigma_c^2 + \Lambda/(N - N_1)} \right] \\ &= \nu^2 \mathbb{E}_{N_1} \left[\sum_{c=2}^{|\mathcal{C}|} \frac{(N - N_1)\tilde{p}_c h(c)}{\nu^2/\sigma_c^2 + \Lambda/(N - N_1)} \right] \\ &= \left(\frac{\nu^2}{1 - p_1} \right) \sum_{c=2}^{|\mathcal{C}|} p_c h(c) \mathbb{E}_{N_1} \left[\left(\frac{(N - N_1)^2}{\nu^2 (N - N_1)/\sigma_c^2 + \Lambda} \right) \right] \end{aligned} \tag{139}$$

Once again, this expression differs from the previous ones in the equal-variance case only as a function

of the target prior variances and $\boldsymbol{N}$. Using a slightly modified version of Lemma 3, we get

$$\begin{aligned} J_T(\boldsymbol{\lambda}^{lo}) &\geq \left(\frac{\nu^2}{1-p_1}\right) \sum_{c=2}^{|\mathcal{C}|} \left(\frac{p_c h(c) N^2 (1-p_1)^2}{\nu^2 N(1-p_1)/\sigma_c^2 + \Lambda}\right) \\ &= \nu^2 N^2 (1-p_1) \sum_{c=2}^{|\mathcal{C}|} \frac{p_c h(c)}{\Lambda + N(1-p_1)\nu^2/\sigma_c^2}. \end{aligned} \tag{140}$$

Similarly,

$$\begin{aligned} J_T(\boldsymbol{\lambda}^{u}) &= \nu^2 \mathbb{E}_{\boldsymbol{C}} \left[\sum_{i=1}^{N} \frac{h(C_i)}{\nu^2/\sigma_{C_i}^2 + \Lambda/N} \right] \\ &= \nu^2 \mathbb{E}_{\boldsymbol{N}} \left[\sum_{c=2}^{|\mathcal{C}|} \frac{N_c h(c)}{\nu^2/\sigma_c^2 + \Lambda/N} \right] \\ &= \nu^2 \sum_{c=2}^{|\mathcal{C}|} \frac{p_c h(c)}{\nu^2/\sigma_c^2 + \Lambda/N} \\ &= (\nu^2 N) \sum_{c=2}^{|\mathcal{C}|} \frac{p_c h(c)}{N\nu^2/\sigma_c^2 + \Lambda} \end{aligned} \tag{141}$$

Note that in all cases, when $\sigma_c^2 = \sigma_0^2$, these results simplify to the expressions under the equal-variance assumption.

APPENDIX XI
DERIVATION OF ASYMPTOTIC BOUND ON MISCLASSIFICATION ERROR

In the regime where $\Lambda \to \infty$, we have a simple expression for the misclassification error using the maximum a posteriori (MAP) estimator. Here, we present the bound for $|\mathcal{C}| = 3$ (i.e., 2 non-zero classes), though the bound is easily extended. The problem then reduces to determining decision regions for when to classify $C_i = 2$ versus $C_i = 3$. From (2), we have

$$f(y_i(t)|C_i = 2) \sim \text{Normal}(\mu_2, \sigma_0^2) \tag{142}$$

$$f(y_i(t)|C_i = 3) \sim \text{Normal}(\mu_3, \sigma_0^2) \tag{143}$$

where we assume that $\mu_2 > \mu_3$ and $p_2 > p_3$ (i.e., $C_i = 2$ is more likely and easier to detect than $C_i = 3$). By the MAP estimator, we have

$$\hat{C}_i = \begin{cases} 2, & f(y_i(t)|C_i = 2)p_2 \geq f(y_i(t)|C_i = 3)p_3, \\ 3, & f(y_i(t)|C_i = 2)p_2 < f(y_i(t)|C_i = 3)p_3 \end{cases} \tag{144}$$

Note that these in terms of the log-likelihood ratio

$$\begin{aligned}\log(LRT) &= \log(f(y_i(t)|C_i=2)p_2) - \log(f(y_i(t)|C_i=3)p_3) \\ &= \log(p_2/p_3) - \frac{1}{2\sigma_0^2}\left((y_i(t)-\mu_2)^2 - (y_i(t)-\mu_3)^2\right) \\ &= \log(p_2/p_3) - \frac{1}{2\sigma_0^2}\left(2y_i(t)(\mu_3-\mu_2) + (\mu_2^2-\mu_3^2)\right)\end{aligned} \tag{145}$$

and

$$\hat{C}_i = \begin{cases} 2, & \log(LRT) \geq 0, \\ 3, & \log(LRT) < 0, \end{cases} \tag{146}$$

Then we have

$$\log(LRT) \geq 0 \Leftrightarrow y_i(t) \geq y_0 \stackrel{\triangle}{=} \frac{(\mu_2^2-\mu_3^3) - 2\sigma_0^2\log(p_2/p_3)}{2(\mu_2-\mu_3)} \tag{147}$$

Moreover, the classification error of Class 3 is given by

$$\Pr(\hat{C}_i = 2|C_i = 3) = \Pr(y_i(t) \geq y_0|C_i = 3), \tag{148}$$

which is simply evaluated using (143). Similarly, the misclassification error for class 2 is given by

$$\Pr(\hat{C}_i = 3|C_i = 2) = \Pr(y_i(t) < y_0|C_i = 2), \tag{149}$$

which is simply evaluated using (142).

REFERENCES

[1] R. Castro, R. Willet, and R. Nowak, "Coarse-to-fine manifold learning," in *Proceedings 2004 International Conference on Acoustics, Speech and Signal Processing*, vol. 3, May 2004, pp. iii–992–5.

[2] R. Castro, R. Willett, and R. Nowak, "Faster rates in regression via active learning," in *Proceedings of the Neural Information Processing Systems Conference (NIPS) 2005*, Vancouver, Canada, December 2005.

[3] E. Bashan, R. Raich, and A. Hero, "Optimal Two-Stage Search for Sparse Targets Using Convex Criteria," *IEEE Transactions on Signal Processing*, vol. 56, pp. 5389–5402, 2008.

[4] E. Bashan, G. Newstadt, and A. Hero, "Two-stage multiscale search for sparse targets," *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2331–2341, 2011.

[5] D. Hitchings and D. A. Castanon, "Adaptive sensing for search with continuous actions and observations," in *Decision and Control (CDC), 2010 49th IEEE Conference on*, 2010, pp. 7443–7448.

[6] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *Information Theory, IEEE Transactions on*, vol. 57, no. 9, pp. 6222–6235, 2011.

[7] J. Haupt, R. Baraniuk, R. Castro, and R. Nowak, "Sequentially designed compressed sensing," in *Statistical Signal Processing Workshop (SSP), 2012 IEEE*. IEEE, 2012, pp. 401–404.

[8] D. Wei and A. Hero, "Multistage adaptive estimation of sparse signals," *Selected Topics in Signal Processing, IEEE Journal of*, vol. PP, no. 99, pp. 1–1, 2013.

[9] M. L. Malloy and R. Nowak, “Sequential testing for sparse recovery,” *arXiv preprint arXiv:1212.1801*, 2012.

[10] M. Malloy and R. D. Nowak, “Near-optimal adaptive compressed sensing,” *CoRR*, vol. abs/1306.6239, 2013.

[11] A. Averbuch, S. Dekel, and S. Deutsch, “Adaptive compressed image sensing using dictionaries,” *SIAM Journal on Imaging Sciences*, vol. 5, no. 1, pp. 57–89, 2012. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/110820579

[12] P. Indyk, E. Price, and D. P. Woodruff, “On the power of adaptivity in sparse recovery,” *Foundations of Computer Science, IEEE Annual Symposium on*, vol. 0, pp. 285–294, 2011.

[13] A. Tajer, R. Castro, and X. Wang, “Adaptive sensing of congested spectrum bands,” *Information Theory, IEEE Transactions on*, vol. 58, no. 9, pp. 6110–6125, 2012.

[14] D. Wei and A. O. Hero, III, “Performance guarantees for adaptive estimation of sparse signals,” Nov. 2013, arXiv:1311.6360.

[15] B. Mu, G. Chowdhary, and J. P. How, “Value-of-information aware active task assignment,” in *AIAA Guidance Navigation and Control*, Boston, August 2013. [Online]. Available: http://arc.aiaa.org/doi/pdf/10.2514/6.2013-4997

[16] L. Bertuccelli and J. How, “Active exploration in robust unmanned vehicle task assignment,” *Journal of Aerospace Computing, Information, and Communication*, vol. 8, pp. 250–268, 2011. [Online]. Available: http://dx.doi.org/10.2514/1.50671

[17] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *J. Amer. Stat. Assoc.*, vol. 58, no. 301, pp. 13–30, Mar. 1963.

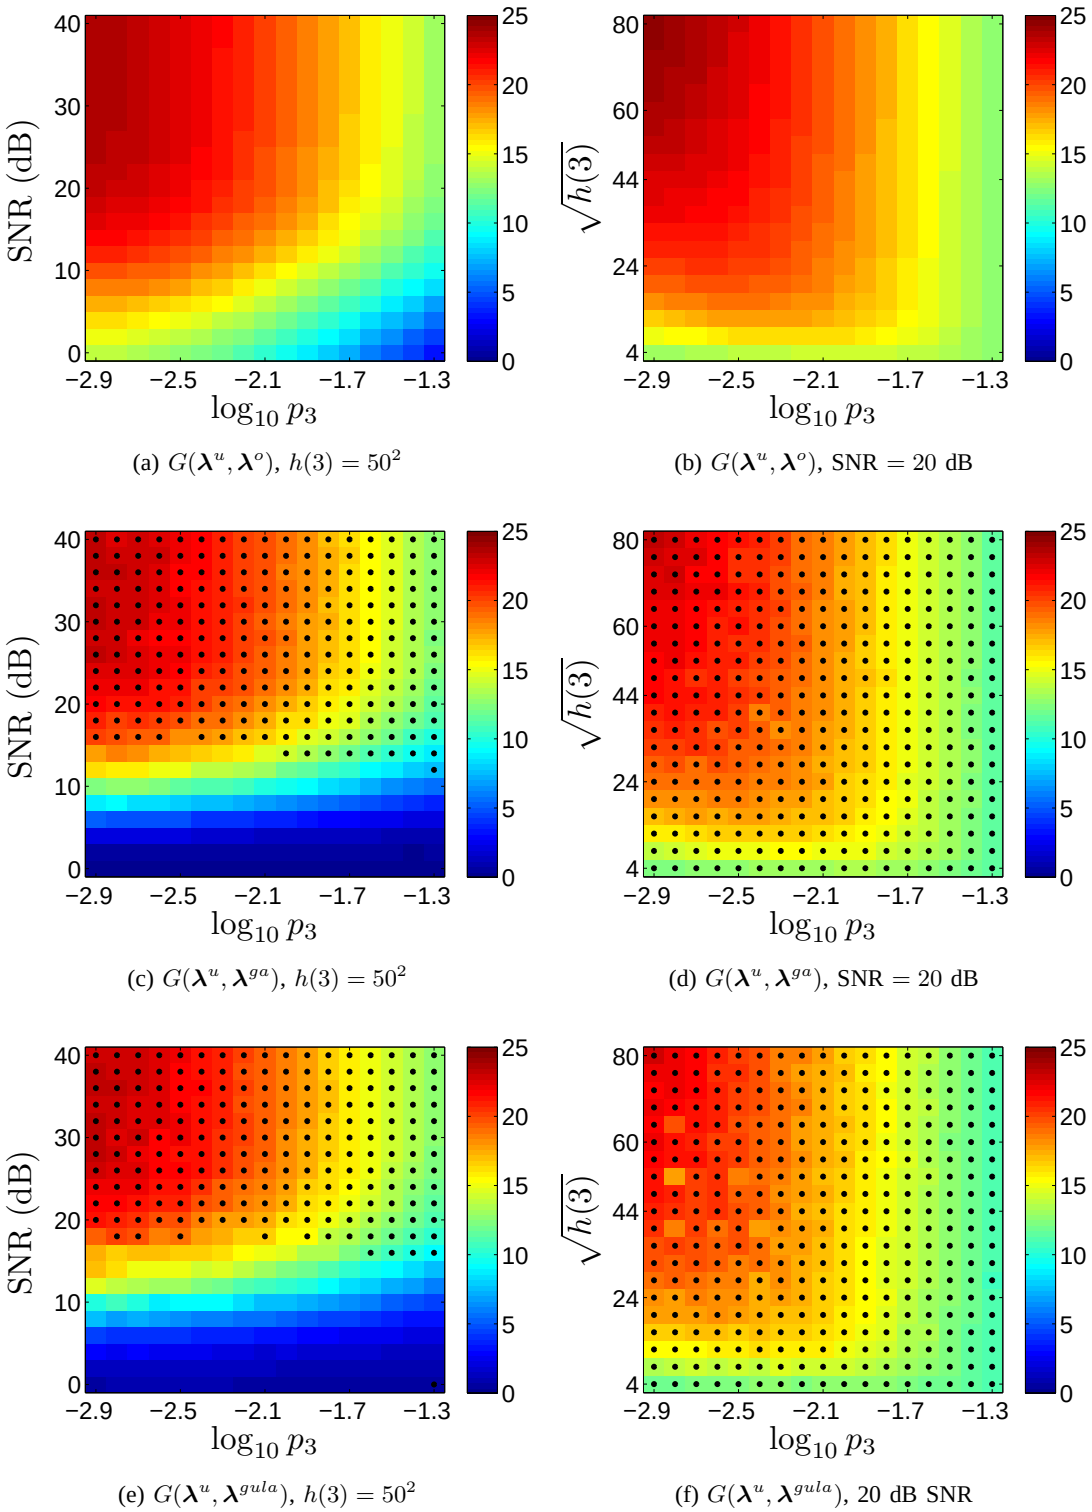

(a) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^o)$, $h(3) = 50^2$

(b) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^o)$, SNR $= 20$ dB

(c) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{ga})$, $h(3) = 50^2$

(d) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{ga})$, SNR $= 20$ dB

(e) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{gula})$, $h(3) = 50^2$

(f) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{gula})$, 20 dB SNR

Fig. 4. Performance gain (in dB) of the GA, GU/LA and full-oracle policies with respect to the GU policy. Gains are shown for varying SNR, class importance weight $h(3)$ and prior probability p_3, with remaining parameters given in Table I. (a) and (b) show results for the oracle policy, (c) and (d) for the GA policy, and (e) and (f) for the GU/LA policy. Black dots indicate that the GA and GU/LA policies are within 3 dB of the oracle policy. For SNR greater than 15 dB, the GA policy achieves within 3 dB of the oracle performance for all values of $h(3)$ and p_3. Similarly, the GU/LA policy comes within 3 dB of the oracle when SNR is greater than 20 dB.

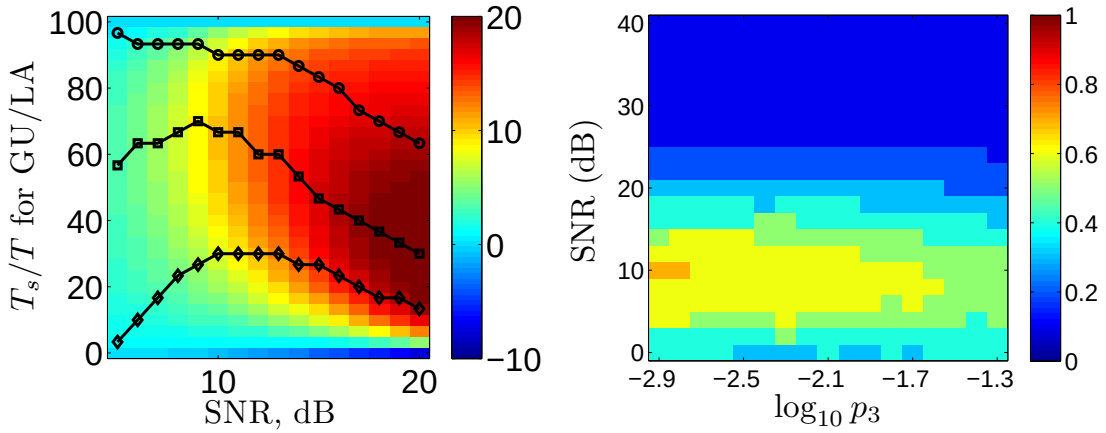

(a) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{gula})$ vs. T_s and SNR

(b) Optimal T_s vs. SNR and p_3

Fig. 5. Sensitivity of GU/LA to T_s/T. In (a), the gain (in dB) of the GU/LA policy is shown as a function of SNR and the percentage of resources used for the global sensor, T_s/T. For each SNR value, the optimal percentage as found by Algorithm 4 is given by black squares. Circles and diamonds indicate the maximum/minimum T_s/T, respectively, where the gain is within 3 dB of the maximum. Note that there is significant degradation when $T_s = 0$ (i.e. the LA policy) or $T_s = T$ (i.e., the GU policy). (b) shows the optimal T_s/T for the GU/LA policy while varying both SNR and p_3.

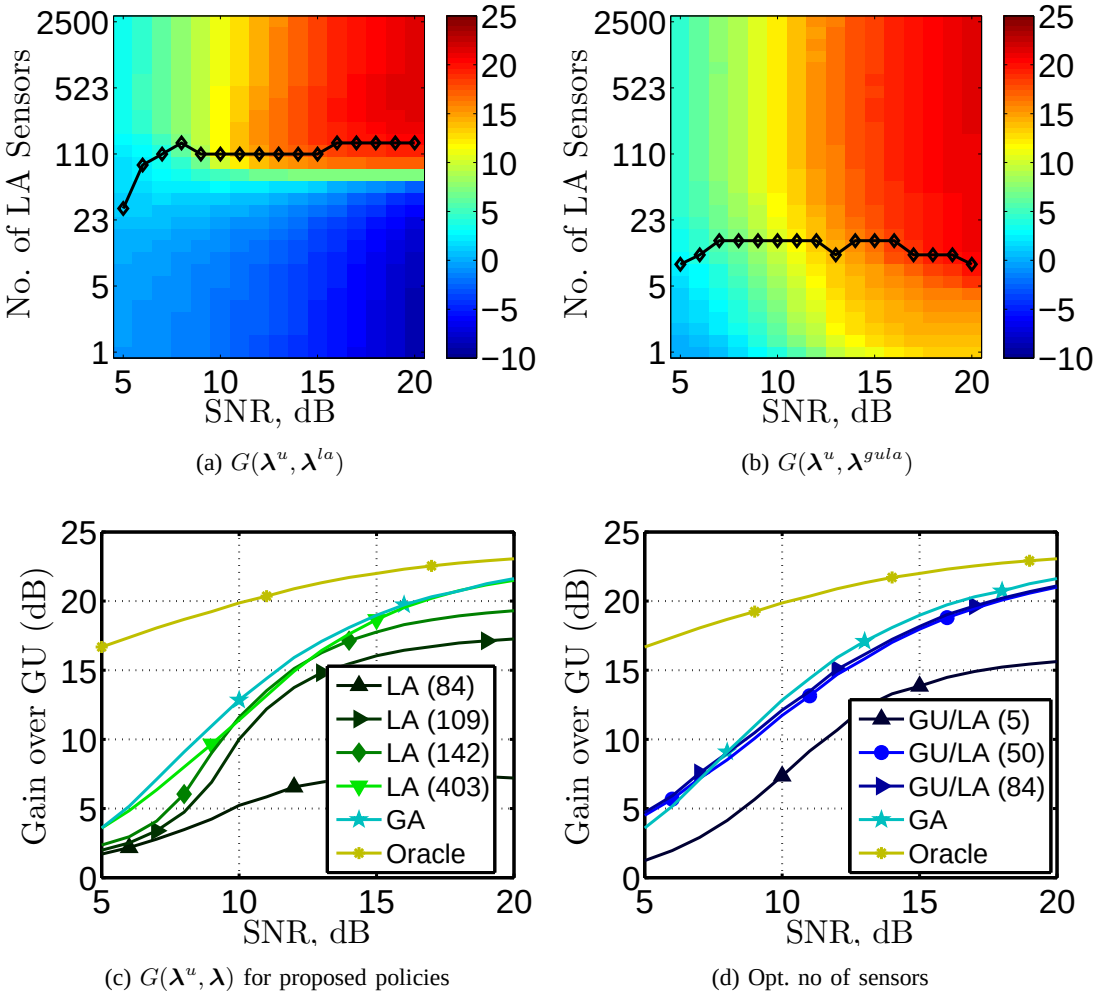

(a) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{la})$

(b) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda}^{gula})$

(c) $G(\boldsymbol{\lambda}^u, \boldsymbol{\lambda})$ for proposed policies

(d) Opt. no of sensors

Fig. 6. Performance benefit by including a global sensor. (a) and (b) show the gains in dB of the LA and GU/LA policies, respectively, over the GU policy as a function of SNR and the number of local sensors. For each SNR value, a black diamond indicates the smallest number of sensors needed to achieve gains within 3 dB of the maximum gain. Both policies have very good performance when the number of local sensors is large. However, the LA policy requires at least 100 sensors to acheive within-3dB performance in most cases, and suffers large decreases in performance when this condition is not satisfied. On the other hand, the GU/LA policy requires many fewer sensors (on the order of 10 sensors) to achieve within-3dB performance. In (c) and (d), the performance of the LA and GU/LA policies are plotted in comparison to the GA and oracle policies. In (c), at least 109 sensors are required to approximate the GA policy. In (d), only 50 sensors are required, while using 84 sensors only provides marginal improvements.

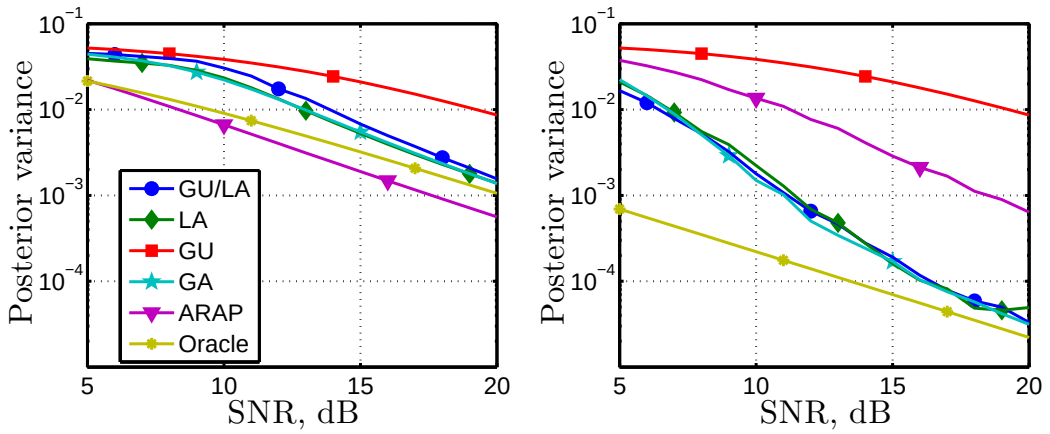

(a) Class 2 targets (b) Class 3 targets

Fig. 7. Posterior variance of the considered policies for both low-importance (class 2) and high-importance (class 3) targets. GU performs the worst as it does not adapt resource allocation. The GA, LA, and GU/LA policies perform better than ARAP over high-importance targets, but worse on low-importance targets. Note that the LA policy uses 400 local sensors, while the GU/LA policy only uses 50 sensors in addition to the GU sensor.

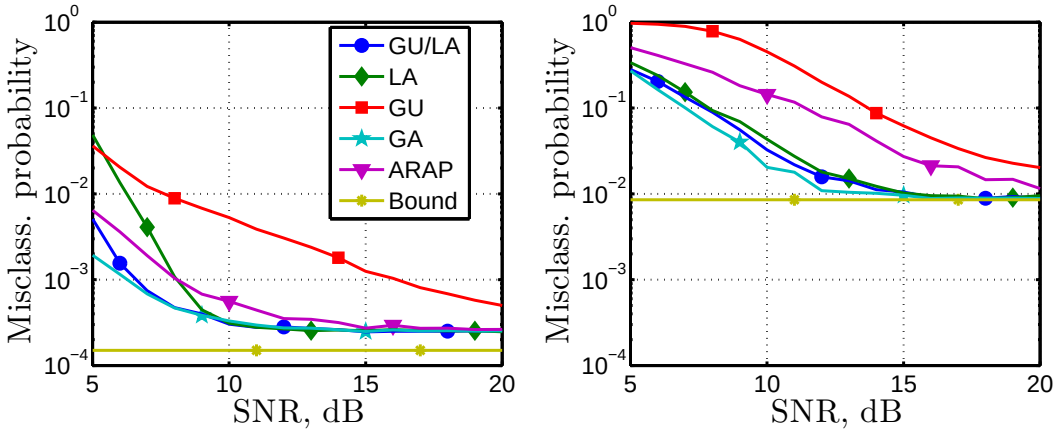

(a) Class 2 targets (b) Class 3 targets

Fig. 8. This figure compares misclassification probability as function of SNR and policy. All adaptive policies perform significantly better than GU and approach the performance of the oracle as SNR gets large. Once again LA, GU/LA, and GA perform better than ARAP in reducing misclassification errors for the high-importance targets.

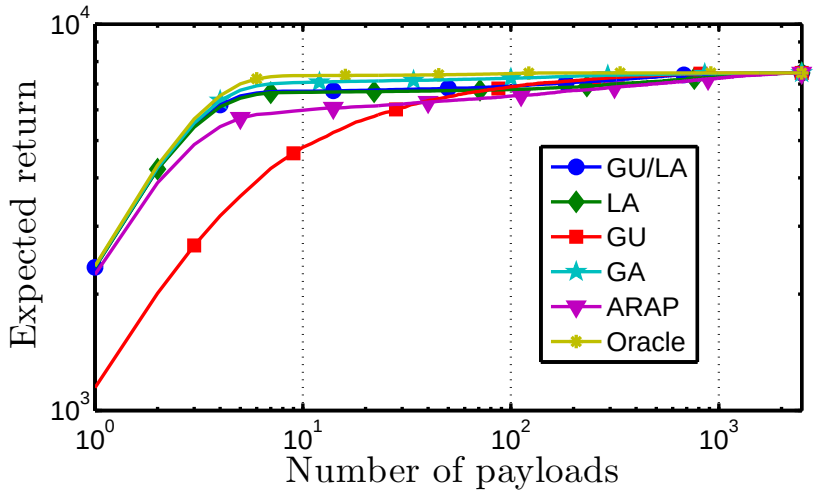

Fig. 9. The expected return given by (54) as function of policy, when there are a limited number of payloads. The GA policy quickly approaches the oracle policy, followed by the GU/LA and LA policies. The ARAP and GU policies have slower convergence to the oracle policy return.